

Historic, Archive Document

Do not assume content reflects current scientific knowledge, policies, or practices.

United States
Department of
Agriculture

Office of
Information
Resources
Management



**National Computer
Center**

Telecommunications
and Applications
Services Center

Version 1.0

August 1993

Guide to the USDA Internet



USDA, National Agricultural Library
NAL Bldg
10301 Baltimore Blvd
Beltsville, MD 20705-2351



United States
Department of
Agriculture

Office of Information
Resources Management
National Computer Center

Telecommunications
and Applications
Services Center

3825 E. Mulberry St.
Fort Collins,
Colorado 80524

March 12, 1994

aTK5105
.875
.I 57G 84
1993

SUBJECT: KA9Q NOS Copyright 1992

TO: KA9Q Users

On Thursday, March 3, 1994, we learned that the copyright for the KA9Q package we have been distributing was not up to date and may have been in error. The most recent copyright for KA9Q dated 1992 states, "Use of this software by commercial or Governmental organizations is on a shareware -- not freeware -- basis." Attached is a copy of the KA9Q copyright.

Because of this new information, we have suspended distribution of the KA9Q software package. In addition, we are advising all users who have received the KA9Q package to either stop using the software or register it if they would like to continue to use it. The registration information is contained in the attached copyright.

We have not suspended the asynchronous dial-in service, only the KA9Q software distribution. Users may continue to use the dial-in service with their own software packages.

Work has begun to evaluate alternative software solutions for the dial-in service. As more information becomes available, we will pass it along. If you have any questions or need assistance with the service, please contact Dan Crosson at (303) 498-1549.

MICHAEL L. THOMAS
Chief
Telecommunications Services Division



OIRM is a Staff
Office of USDA

AN EQUAL OPPORTUNITY EMPLOYER



KA9Q NOS Copyright 1992

The "KA9Q NOS" networking package is copyright 1992 by Phil Karn, KA9Q. All rights reserved. Permission is hereby granted for the free copying and use of this package by radio amateurs and non-profit educational institutions. I.e., to hams and schools, it is "freeware".

Use of this software by commercial or governmental organizations is on a "shareware" -- not freeware -- basis. That is, a prospective commercial or governmental user may copy and evaluate the software for his or her application. If it is to be used beyond evaluation purposes, it must be registered for \$50/copy (see address below).


The right conferred by the registration fee to use the software in a commercial or government organization is limited to those portions of the NOS package that are the sole work of Phil Karn, KA9Q. Certain portions of the package, in particular the email subsystem, were contributed by others with the understanding that they would be used for amateur radio and educational purposes only. Any commercial or governmental use of these portions requires an agreement with the appropriate author(s). Each source code file includes an appropriate authorship and copyright statement.

In several cases, the contributing authors have placed their code into the public domain. This is the case with the PPP and VJ header compression code contributed by Katie Stevens and Bill Simpson. This code may be freely used by anyone. Again, see the comments in the various source files for details.

NOS is provided without warranty of any kind. The prospective user is responsible for testing and evaluating it to see if it meets his or her needs. The payment of the registration fee merely grants the user the right to use the software as-is; it does not obligate me to provide any assistance or support.

OEM licenses for the inclusion of part or all of the Karn-authored portions of NOS into commercial networking software products are available on an individually negotiated basis. Discounted site licenses (e.g., for large-scale internal use) are also available. Please contact:

Phil Karn, KA9Q
7431 Teasdale Ave
San Diego, CA 92122
Voice: (619) 587-8281
Fax: (619) 587-1825



United States
Department of
Agriculture

Office of
Information
Resources
Management


National Computer
Center

Telecommunications
and Applications
Services Center

Version 1.0

August 1993

Guide to the USDA Internet



Contributors, Authors, Editors and Publishers

Telecommunications Services Division

Jack Barner
Dan Crosson
Gordon Durflinger
Ike Floyd
Doug Harsha
Truman Harsha

Telecommunications Policy Division

Diana Carriker
Hampton Hoge

Agency Applications Services Division

Mary Mosier

Washington Service Center

Craig Leavitt
Clif Gonzales
Ron Hobbs
Terry Scalsky

National Computer Center

Gary Laffoon
Greg Schmitz

Computer Data Systems, Inc.

J. Bert Noel



Table of Contents

Log of Modifications	v
Preface	vii
1. Introduction	1
2. Scope	3
3. Background	5
4. References/Abbreviations/Definitions	7
4.1 References	7
4.2 Abbreviations	7
4.3 Definitions	8
4.4 Mailing Addresses / Telephone Numbers / Internet Addresses	8
5. Policy Information	11
6. Description of Facilities	13
6.1 Overview	13
6.2 Facilities	14
Fort Collins Internet Gateway	14
Washington, DC Concentrator	15
Kansas City Concentrator	16
FmHA Network	17
Network Resources	17
7. Description of Access Methods	19
7.1 Dedicated Access	19
7.2 Asynchronous Access	19
SLIP	19
Terminal Emulation	20
7.3 X.25	20
8. Services	21
8.1 Administration Services	21
Requesting Services	21
IP Address Allocation	21
Description of IP Addresses	21
Requesting IP Addresses	22
Domain Name Allocation	23
Description of Domain Names	23
Requesting Domain Names	24

Request to Manage a Domain	24
KA9Q Distribution	25
Description of KA9Q Software	25
Requesting KA9Q Software	25
Installing the KA9Q Software	25
Using the KA9Q Software	26
8.2 Network Services	26
Common Services	26
Domain Name Server	26
Description of the Domain Name Server	27
Using the Domain Name System	27
Help Desk	27
Description of the Help Desk	28
Requesting Help Desk Service	28
Guest Host Services	29
Mailbox	29
Description of Mailbox Service	29
Requesting Mailbox Service	29
Using Mailbox Service	29
Telnet	30
Description of Telnet	30
Requesting Telnet Service	30
Using Telnet	30
Telnet Documentation	31
Other TCP/IP Applications	31
Description of Other TCP/IP Applications	31
Requesting Other TCP/IP Applications Service	31
Using Other TCP/IP Applications	32
FTP Documentation	32
8.3 Router Configuration Service	32
Description of Router Configuration Service	32
Requesting Router Configuration Service	33
8.4 Engineering Services	33
Description of Engineering Services	33
Requesting Engineering Services	33
9. Cost Recovery	35
10. USDA Resource Center Usage	37
10.1 NCC Applications	37
10.2 Other USDA Resource Centers	37

Appendices

APPENDIX A

USDA Internet Asynchronous Access Request Form	39
--	----

APPENDIX B

USDA Internet Guest Host Services Request Form	41
--	----

APPENDIX C

IP Subnetwork Address Request Form	43
--	----

APPENDIX D

USDA Domain Name Request Form	45
-------------------------------------	----

APPENDIX E

USDA Domain Management Request Form	47
---	----

APPENDIX F

KA9Q User Reference Manual Supplement	49
---	----

APPENDIX G

Learning More About the Internet	57
--	----

APPENDIX H

Zen and the Art of the Internet	63
---------------------------------------	----

Log of Modifications

This section will summarize the changes made to this document as it evolves. Since this is the first version of this document, there are no changes listed.

<u>DATE</u>	<u>Changes Made</u>
-------------	---------------------

Preface

During the departmental review of this document a number of comments were made concerning its name, or more specifically, the phrase "USDA Internet". Suggestions ranged from using the phrase "USDA TCP/IP Network" to simply the "USDA Network" or still simpler the USDANet.

The USDA Internet is comprised of an interconnection of networks "owned" and operated by the USDA. It is certainly not the same as "The Internet"; however, it relies upon "The Internet" to interconnect with many other important networks. The USDA Internet uses FTS2000 circuits and The Internet does not use FTS2000 facilities. This is an important distinction since FTS2000 mandates that USDA use network A circuits whenever intra-USDA business is conducted across Local Access and Transport Areas (LATA).

The name that is chosen for this document is certainly important since the intuitive value of a name can determine success for the product it represents. It needs to convey that this is not a public network. The name should certainly imply the protocols which can be used on that network. The title of this document, "Guide to the USDA Internet" could have avoided this controversy with a title like, "Guide to USDA Internetworking" or "Guide to the AGNET" or something else which obscures the contention. What the resources that make up this network are eventually named is not the critical issue at this time. Rather, it is imperative that USDA planners and networkers understand that there is something tangible out there that will support internetworking on a national scale with proven, commercial off the shelf software and hardware.

Douglas Comer in his text, "Understanding TCP/IP" goes to considerable effort to explain the term "Connected Internet". He refers to the high speed backbones, such as NSFNet with this term to distinguish those networks which comprise the core of the Internet from networks such as USDA's which is a network restricted to USDA users. While IBM or DEC have a substantial Internet involvement, for example, they are certainly not a part of the Internet backbone even though they have many interconnected networks distributed throughout the world.

In fairness to the readers and the many contributors/reviewers of this document, a list of the names not used for this networking enterprise are enumerated along with the reason for not choosing them at this time.

- USDA Network - (too broad, what about SNA nets, DEC nets, X.25 PSNets, et. al.)
- USDANet - (same as above, implies the Internet backbone like NSFNet)
- USDA TCP/IP Network - (too long and too specific. Routers on this network are going to be GOSIP conformant implying the possibility for OSI protocols)

USDA Internetwork and AGNet were also considered.

In any case, the title and text will continue to use the phrase, "USDA Internet" for the time being. This document represents a starting point. It will undoubtedly evolve, possibly into more than one document. Suggestions for improvement are welcomed and encouraged.

1. Introduction

The purpose of the *Guide to the USDA Internet* is to assist in the planning, designing, implementation, operation and maintenance of an Internet connection using the USDA Internet. A national help desk is being established to provide assistance to all USDA agencies in establishing Internet connections. This document provides descriptions of the facilities and services offered as well as important information regarding the use and the cost of the services. Several appendices are included with this document. They include several forms which may be used to request USDA Internet services, and other documents which provide more detailed background on the organization and operation of the Internet. Appendix G includes a short description on how a new Internet user can learn more about the Internet with the application called "Gopher." Appendix H is a reprint of a paper: *Zen and the Art of the Internet*, by Brendan P. Koehoe of Widener University. It provides an excellent introduction to the workings and "culture" of the Internet. It is included in this document by kind permission of the author.

The use of the Internet in the USDA is expected to continue to expand. This Guide will be updated as new features and information become available. The OIRM Telecommunications Services Division (TSD) welcomes comments and suggestions about this Guide. Information on how to contact TSD is included later in this document.

This Guide will probably be of greatest value to USDA agency network administrators and other computer professionals who need access to the Internet. There is also information included in this document that will benefit general users. This document's principal focus is not that of a tutorial, although there are several suggestions on finding more information about the Internet and the protocols and applications which are used on it.

2. Scope

The scope of this document includes all USDA agencies that choose to participate in the USDA Internet. This Guide to the USDA Internet describes the USDA Internet and the general services and capabilities that are being offered to all USDA agencies that have a requirement to access the Internet.

This Guide is intended for a very broad audience ranging from network administrators to relatively inexperienced PC users. This will undoubtedly cause some readers to say it is too technical while others will want more detail. During review of this document there were a number of comments that the document should either be reorganized or split up into multiple documents. It has always been the intention of the USDA Internet team to provide follow on information either through other printed matter or through some on-line resource like a Gopher server. This document represents a good start to a large and complex task. It will probably go through a number of revisions during its lifetime. Readers are encouraged to provide input on their requirements and suggestions for making this guide better.

3. Background

The Internet is a world-wide collection of networks that support commercial, educational, and governmental user communities. This network has evolved from an experimental network in the 1970's, into a massive operational network that provides connectivity for a vast user community. It is by far the single largest data highway system in the world and access to the Internet has become extremely important to many USDA agencies during the past few years.

Because of the significant benefits associated with Internet access, several USDA agencies have taken the initiative to develop their own access facilities. These individual agency efforts have filled an important need and should be recognized for their leadership and initiative. Even so, the costs associated with the continued development and implementation of facilities with duplicate functions have prompted OIRM to design, develop and implement a departmental Internet access facility.

The diverse range of USDA requirements and capabilities to access such a network has lead OIRM to design and develop a general set of services and access methods. Additional plans have been made to enhance the connectivity options, the location of major concentration nodes and access methods. The objective of this capability is to provide a cost effective, high performance access to this extremely important network resource. An associated benefit is that USDA nodes participating in the network will be able to share data easily.

The USDA Internet provides a network medium upon which many applications will be able to operate. The most well known of these applications are Telnet, FTP and SMTP. Telnet allows a user to sign on to another computer or host. FTP stands for File Transfer Protocol and allows a user to either send or retrieve files to or from another host. SMTP stands for Simple Mail Transfer Protocol and provides a facility for users to send messages to other users on the Internet.

4. References/Abbreviations/Definitions

4.1 References

Following are references that provide additional insight to the topics addressed in this Guide. RFC stands for Request for Comment. This is the standard method for organizing documentation about the Internet.

RFC 920; Postel, J.B., Reynolds, J.K. *Domain Requirements*, October 1984.

RFC 1032; Stahl, M.K. *Domain Administrators Guide*, November 1987.

RFC 1034; Mockapetris, P.V. *Domain Names - Concepts and Facilities*, November 1987.

RFC 1035; Mockapetris, P.V. *Domain Names - Implementation and Specifications*, November 1987.

RFC 1338; Fuller, V., Li, T., Yu, J.; Varadhan, K. *Supernetting: an Address Assignment and Aggregation Strategy*, June 1992.

RFC1118; Krol, Ed. *The Hitchhikers Guide to the Internet*, September 1989.

Comer, Douglas E., *Internetworking with TCP/IP*, Volume 1, 1991.

4.2 Abbreviations

Following is a partial list of abbreviations that are used in this user guide.

DAR	Designated Agency Representative
DNS	Domain Name System
DTS	Dedicated Transmission Service
FTP	File Transfer Protocol
IP	Internet Protocol
KCMO	Agriculture Stabilization and Conservation Service's (ASCS) Kansas City Management Office
LAN	Local Area Network
NCC	National Computer Center, Kansas City, MO
NFC	National Finance Center, New Orleans, LA
NGC	The Soil Conservation Service (SCS) National Cartography and Geographical Information Systems Center (Ft. Worth)
NTC	a Soil Conservation Service National Technical Center
OSI	Open System Interconnection
PC	Personal Computer
PSS	Packet Switched Service
RFC	Request For Comment
SLIP	Serial Line Internet Protocol
SMTP	Simple Mail Transfer Protocol

T1	a high-speed communication circuit equivalent to 24 56-Kbps circuits
TASC	Telecommunications and Applications Services Center
TCP/IP	Transmission Control Protocol/Internet Protocol
TCP	Transmission Control Protocol
TELNET	Standard Internet terminal emulation protocol
TSD	Telecommunications Services Division, Fort Collins, CO
WSC	Washington Service Center, Washington, DC

4.3 Definitions

Connected Internet	Term used to refer to the world's largest internetwork, connecting thousands of networks worldwide.
IP address	32 bit Internet address assigned to hosts or devices using TCP/IP.
Router	An OSI Layer 3 device used to connect networks that can determine which of several paths network traffic will follow.
USDA Internet	Term used to refer to the network of USDA agencies interconnected to each other and to the connected Internet.

4.4 Mailing Addresses / Telephone Numbers / Internet Addresses

OIRM Telecommunications Services Center
 3825 E. Mulberry St.
 Fort Collins, CO 80524
 303/498-2007
 FAX 303/498-1660
 addman@ag.gov or A76ADDR@ATTMAIL.COM

Washington Service Center
 Internet TCP/IP Administrator
 USDA-OIRM-WSC
 14th and Independence Ave. SW
 Washington, DC 20250
 Room South -162
 202/720-6828
 FAX 202/720-8274
 admin@wsc.ag.gov

National Computer Center, Kansas City
Customer Support Center
USDA-OIRM-NCC
8930 Ward Parkway
Kansas City, MO 64114
816/926-6681
FAX 816/926-6754
a76hvaughn@attmail.com

5. Policy Information

Departmental Regulation (DR) 3300-1, Telecommunications, Appendix I draft, provides policy guidelines for USDA agency participation in the Internet. It includes statements of policy covering the definition, allocation, and management of Internet addresses; the definition, allocation, and management of the Domain Name System (DNS); and defines rules for proper use of the resources of the Internet by USDA agencies.


At this writing USDA policy concerning the use of Internet has not been finalized. When it has been finalized, it will be located in Departmental Regulation (DR) 3300-1, Appendix I and will not be included as part of this document.

While procedures for obtaining resources are described in another section, the reader is reminded that agency-specific policies may prescribe that only certain designated representatives may be allowed to actually submit requests for those resources to OIRM. Users are encouraged to consult your agency's policy for guidance in this area.

6. Description of Facilities

6.1 Overview

The facilities being implemented to support generalized USDA Internet access are designed to offer multiple access methods as well as a wide range of services. Initially, major nodes will be located in Fort Collins, Colorado and Washington, DC. Another major node in Kansas City, Missouri is planned for the future. Additional nodes will be connected as agency requirements are indicated. The following sections describe the specific configurations of hardware and software that are located at each major node. Figure 6.1 illustrates the initial deployment of the USDA Internet.



6.2 Facilities

Fort Collins Internet Gateway

The Fort Collins node is configured to support high speed (i.e. T1) access to the Internet via Colorado SuperNet. Remote USDA customers have access to this node via three telecommunications access methods: dedicated circuit access, X.25 access and asynchronous access. Refer to Section 7 for a detailed description of access to the USDA Internet. The Fort Collins Internet Gateway facility is illustrated in Figure 6.2.




Figure 6.2

Washington, DC Concentrator

The Washington, DC node is configured to support access to the Internet via the Fort Collins Internet gateway. Remote USDA customers have access to this node via three telecommunications access methods: dedicated circuit access, X.25 access and asynchronous access. Refer to Section 7 for a detailed description of access to the USDA Internet. The Washington, DC concentrator facility is illustrated in Figure 6.3.




Figure 6.3

Kansas City Concentrator

As of June 1993, The National Computer Center (NCC) located in Kansas City, Missouri was in the process of implementing Transmission Control Protocol/Internet Protocol (TCP/IP) on its IBM MVS Host Complex. Initially, an FmHA router will be used to provide connectivity to the USDA Internet. A permanent replacement is expected to be installed in the future which will triangulate the central USDA Internet backbone and provide connectivity to services located at the NCC.

Procurement is currently underway to provide channel attached connectivity from the IBM host to the local IP router. TCP/IP access to the NCC will be available using the FTS2000 X.25 Packet Switched Service and a dedicated circuit to connect the NCC router to the backbone network. The NCC Kansas City concentrator facility is illustrated in Figure 6.4.

NCC Kansas City Concentrator




Figure 6.4

FmHA Network

FmHA has routers in St. Louis, Kansas City, and Washington, D.C. One of the St. Louis routers is connected to an FmHA router at NCC-KC; the FTS2000 Packet Switch Network; an FmHA router at the South Building in Washington, D.C.; and other routers in St. Louis. It is also connected through the Internet via a circuit to George Washington University in St. Louis. These network connections utilize 56Kbps circuits, except for four routers in St. Louis which are connected via FDDI. At the time of this writing, the FmHA Washington router is not connected to the South Building Ungeremann-Bass LAN. The FmHA router at NCC-KC is connected via 56 Kbps DTS FTS2000 circuit to the backbone router in Fort Collins and via X.25 to the mainframes at NCC-KC. There are plans to connect the router at NCC-KC to a LAN which will have connectivity to the mainframe as well as a connection to the ASCS LAN at KCMO.

Network Resources

Initial network resources will consist of the following:

1. A T1 connection to the Internet via Colorado SuperNet at Colorado State University in Fort Collins, CO.
2. One Cisco Router in Fort Collins and one in Washington, DC initially configured to handle twelve (12) low speed ports (56K), two (2) High Speed Ports (T1), and two (2) ethernet ports.
3. A twenty (20) line Asynchronous Telephone Rotary connected to a communications server in Fort Collins and in Washington, DC for dial up asynchronous access at speeds up to 9600 baud for use with the Serial Line Internet Protocol (SLIP).
4. RISC based hosts located in Fort Collins and in Washington, DC will provide Domain Name Service, password and ID verification for dial-in users, and act as a TCP/IP application platform for non SLIP asynchronous users.

The initial network configuration allows FTS2000 X.25 packet use, dedicated FTS2000 access, and dial-in asynchronous use via SLIP. Access is through either Fort Collins or Washington, DC with future connectivity via Kansas City. All connections made to the gateway router in Fort Collins from either Washington, DC or Kansas City will be made via dedicated FTS2000 DTS 56K lines or X.25 packet switched service. As the network load increases, additional lines will be installed as required to meet demand.

7. Description of Access Methods

The access methods being implemented to support generalized USDA Internet access are designed to offer USDA agencies options at various levels of performance and cost. The three methods currently provided are dedicated, asynchronous and X.25. Each of these access methods is described in the following sections.

7.1 Dedicated Access

Dedicated FTS2000 access is defined as the FTS2000 service known as Dedicated Transmission Service (DTS). This service can be ordered by an agency via their USDA Designated Agency Representative (DAR) if they have a need for higher speed access to the Internet than is provided via dial-in asynchronous service. The agency must be willing to incur the FTS2000 costs associated with installation and maintenance of a DTS line into the router in either Washington, DC or Fort Collins, CO. If there is a need to order an additional router card to terminate this DTS line, a portion of that expense will also be incurred by the USDA agency requesting DTS access. All requests for this type of access will be coordinated through the router configuration manager of the router to which a connection is needed. The appropriate facilities manager can provide the FTS2000 T1 channelization information which the agency will need to place a DTS order.

7.2 Asynchronous Access

There are two asynchronous access methods provided by the USDA Internet. Asynchronous communication is defined as transmitting and receiving data without precise clocking. This is the type of communication that is often used to connect a PC to a host using a modem and a telephone line. This type of communication eliminates the need for a dedicated line. The two types of asynchronous access provided on the USDA Internet are Serial Line Internet Protocol (SLIP) and terminal emulation described in sections below.

SLIP

Serial Line Internet Protocol (SLIP) defines a method of sending Internet packets over standard RS-232 asynchronous serial lines such as the public switched telephone network. In the USDA Internet, SLIP is used to allow PC users access to the Internet via serial lines. SLIP would not be used if a direct Internet connection such as a LAN or mainframe host is available.

The SLIP protocol defines a method of encapsulating IP packets on serial lines. A characteristic of the SLIP access method is that all TCP/IP software and applications (Telnet, FTP etc.) are provided by software running on the PC. Because of this, users using the SLIP access method must have a TCP/IP software package capable of SLIP communication installed on their PCs.

There are many software packages available in the marketplace and in the public domain that offer TCP/IP and SLIP. Any of these packages are acceptable to use when establishing a SLIP connection to the USDA Internet. The public domain KA9Q package is being distributed by

OIRM. A description of how to obtain a copy of the KA9Q software package is included in Section 8 of this Guide.

To access the USDA Internet using SLIP, the user will connect to a terminal server using a SLIP software package on a PC, a modem and a serial line. Once connected the user will sign onto the terminal server using the userid and password supplied by one of the service centers in Fort Collins, Washington, DC, or Kansas City. After signing onto the terminal server the "slip" command is issued to put the communications line into SLIP mode. At this point the PC has a TCP/IP connection to the USDA Internet. The applications provided by the SLIP package installed on the PC may then be used to interact on the Internet.

Terminal Emulation

Terminal emulation access to the USDA Internet is another access method for point-to-point serial connections. This access method is used to connect a terminal or a PC to a remote host over standard RS-232 asynchronous serial lines such as the public switched telephone network. A characteristic of the terminal emulation access method is that all the TCP/IP software and applications are provided by a remote host or communications server directly connected to the USDA Internet.

To use this access method the user will connect to a communications server using some type of terminal emulation software package (Kermit, Crosstalk and Smarterm for example) across a serial line using modems and the public switched telephone network. After the user has established a terminal session, the communications server will provide the standard Internet terminal emulation application which is called Telnet. Telnet may be used to access any host connected to the Internet. To gain access to other TCP/IP applications such as FTP or Mail, the user will need to use Telnet to get access to a host providing those services. Telnet and other host supplied services are discussed in Section 8 of this Guide.

7.3 X.25

FTS2000 X.25 packet access will be available via the Fort Collins router upon completion of an initial testing phase (approximately August 1993). This testing will determine the feasibility of X.25 use for Internet access and the ability to rotary multiple X.25 lines for the most efficient utilization of the packet access. This access will be coordinated with the facility manager of the Fort Collins router.

8. Services

OIRM provides several important services for the users of the USDA Internet. These services include administrative services, network services and engineering services. They are described in the following sections.

8.1 Administration Services

Several important services will be performed by OIRM for the USDA Internet user community. Those services include the allocation of IP subnetwork addresses and domain names and the distribution of KA9Q (PC/SLIP) packages.

Requesting Services

Forms for requesting the services described below are included as appendices to this Guide. Electronic versions of these forms can be obtained by sending a request to:

“A76ADDR@ATTMAIL.COM” or “addman@ag.gov”

To request a service, complete a request form and send the request via your requesting agency's Internet Coordinator to OIRM/TSD Fort Collins, preferably via E-Mail:

“A76ADDR@ATTMAIL.COM” or “addman@ag.gov”

Forms may also be mailed or FAXed to the addresses and numbers listed on each form.

IP Address Allocation

The logical addressing of devices on a network is accomplished through the use of Internet Protocol (IP) addresses. These addresses must be allocated to ensure that their values are unique. A complete inventory of these addresses shall be maintained by the USDA addressing authority, OIRM/TSD.

Description of IP Addresses

IP addresses are 32 bit quantities which uniquely describe any single device. That device could be a super computer or a laser printer. There are several types of IP addresses. OIRM/TSD has received allocations of both class B and C addresses.

Class B addresses use 16 bits to describe the network and 16 bits to describe the host. Class C addresses use 24 bits to describe the network and only 8 bits to describe the host. The host portion of both the class B and C addresses can be further subdivided to implement what is commonly known as subnetting.

The three class B addresses that are assigned to OIRM represent an addressing resource capable of supporting over 180,000 hosts/devices. OIRM has elected to subnet one class B address with 8 bits for subnetworks and 8 bits for hosts/devices for the USDA LAN in Washington, DC. This will allow 254 subnetworks each capable of addressing 254 hosts/devices. The other two class B addresses are subnetted with 9 bits for subnetworks and 7 bits for hosts/devices. This will allow 510 subnetworks each capable of addressing 126 hosts/devices.

The IP network address resource is growing more and more scarce with each passing day. The assignment of IP subnets is difficult with this address constraint because large segments of a particular network address should not be arbitrarily parceled out. The number of devices and the topology of the physical network must be considered if an efficient, functional allocation is to be made. It is possible that arbitrary assignments can lead to network performance problems or even network failures.

Requesting IP Addresses

IP addresses will be allocated, whenever possible, as groups or subnets of IP class B addresses. Most devices requiring IP addresses are normally found on local area networks (LANs). Improperly designed address allocations can cause poor network performance or even network failure. When more than one subnet is required for a particular physical network (e.g., a LAN), those subnets should be allocated as a contiguous block of subnets to maximize performance and minimize configuration management problems. For further reading on subnets see the RFC 1338 (Supernetting).

The form for requesting an IP subnetwork address is reproduced in Appendix C. These requests should include sufficient justification for obtaining a subnet based on the USDA policy documented in Departmental Regulation (DR) 3300-1, Telecommunications, Appendix I. The Washington Service Center is also available to issue Internet addresses for the convenience of Washington, DC-area users. For those agencies that prefer a single OIRM contact for their agency's coordinator, arrangements should be made with OIRM/TSD.

Requests for IP address resources must contain information which is described below. The items included on the form are defined as follows:

Agency Name is the acronym for the particular agency such as SCS.

Unit Name is the name of the agency entity which operates the particular subnetwork.

Location of subnetwork is the physical location of the subnetwork.

Administrative Authority is the individual responsible for this particular subnetwork, usually a manager.

Technical Contact is the individual actually operating and managing the subnetwork.

Description of subnetwork is information which is essential for an efficient and functional address allocation to be made. Make and model of routers, the routing protocols which the router supports, the type of subnetwork, and numbers of devices on the network are all pieces of information which are important to the allocation process and also impact planning for network upgrades.

Is subnetwork connected to the Internet? If a subnetwork is connected to the Connected Internet and not connected to the USDA Internet, then a USDA subnetted class B address may not be used to satisfy the request. In that case a request is created by OIRM and forwarded to the Internet addressing authority, the InterNIC, for one or more class C addresses to satisfy the request. The agency making the request will be notified once the class C allocation has been made. (Note: These requests should conform to USDA policy documented in Departmental Regulation (DR) 3300-1, Telecommunications, Appendix I.)

Domain Name Allocation

Domain names are human-readable strings of characters which uniquely describe entities which can then be accessed via the Internet. OIRM/TSD is the authority for the domain **ag.gov** and is therefore providing the service of domain name allocation for **ag.gov**.

Description of Domain Names

Domain names are more intuitive strings of characters than IP addresses. They are transparently translated by a "name server" into IP addresses. Name servers exist throughout the Internet and are used by applications to translate domain names into IP addresses which are then used by those applications to communicate with their peers on other hosts.

Domain names are assigned hierarchically under various first level domain names, such as "gov" for use by the civilian side of government, "mil" for use by DoD, "com" for use by commercial entities and "edu" for use by universities. The second-level domain name of **ag.gov** has been allocated for the USDA's use by OIRM/TSD, the USDA addressing authority. DR 3300-1, Appendix I will provide for a starter set of third level domain names based on the NFC agency acronyms (fns.ag.gov, for example).

Compelling reasons for the use of other second-level and third-level domain names may exist. DR 3300-1, Appendix I allows for this possibility and provides a procedure for agencies whose needs are not best met by the <agency_acronym>.ag.gov third-level qualifiers.

Domain names must be unique. In other words, a mail host may have a domain name of **mail.ag.gov**. John Doe's address on that mail system might then be **jdoe@mail.ag.gov**. The domain name **mail.ag.gov** must identify one and only one computer. The converse is not true, however. Computers may have more than one legitimate name. That very same mail host might also be known as **mail.oirm.ag.gov**. In either case both domain names

would be linked to the same IP address. For further reading on the Domain Name System see RFCs 1032, 920, and 1034.

Requesting Domain Names

To obtain Domain Names you must complete the request form which is reproduced in this guide as Appendix D. These requests should include sufficient justification for obtaining a domain name based on the USDA policy documented in Departmental Regulation (DR) 3300-1, Telecommunications, Appendix L

The terms used in the form are defined as follows:

Requested Name is the domain name that is being requested. For example Food and Nutrition Service might request a domain name for a mail host of "mail.fns.ag.gov".

IP Address is the IP address of the device named above.

The other terms are defined in the section "Requesting IP Addresses."

Request to Manage a Domain

While OIRM will operate several name servers for the Department, there will be cases where units will want to operate their own name servers and manage their own domains.

Because domain names are hierarchical in nature, they are often described by a tree. For the USDA there will be at least one root domain, ag.gov. Under that root, which is only a small portion of the name-space tree, there will be many third-level qualifiers, or sub-trees. In the event that a particular unit decides to manage its own domain, then a portion of that tree would be "carved" out for that unit to manage. In the case where a USDA agency called "abc" decided to manage its own domain, then the third-level domain name of "abc.ag.gov" could be delegated to a unit of that agency and that unit could in turn issue fourth-level (and so on) names to their agency's other units.

There are certain prerequisite conditions which must be satisfied before such an arrangement can be implemented. Principal among these prerequisites is that a name server be operational and be backed up by one or more secondary name servers which are not located at the same location as the primary. In other words, the primary and secondary name server may not be positioned in such a way that they share a "single point-of-failure." An OIRM name server should be included as one of the secondary name servers.

The form which is used for requesting a domain delegation is reproduced in Appendix E.

KA9Q Distribution

One of the administrative services offered by OIRM to support access to the USDA Internet is the distribution of the KA9Q software package which runs under the DOS operating system. The following sections describe the KA9Q software package and how to request, implement and use the package.

Description of KA9Q Software

The KA9Q software package contains network protocol family commonly known as TCP/IP. It is the result of a developmental effort lead by Phil Karn and is available for general use in the Internet community. It will be one of the packages used in the USDA Internet to provide TCP/IP communications using SLIP across serial connections. Other implementations of TCP/IP using SLIP are also valid and may provide additional functionality, but the KA9Q package has the advantage of being available at no cost.

The software package contains all the software and documentation necessary to install the KA9Q software on a DOS PC and establish a SLIP connection to the Internet. The PC should be connected to the public switched phone network through a modem attached to one of the communications ports on the PC. The next section describes how to obtain the KA9Q package and an account on the terminal server. **Note:** even if another SLIP package is being used, the user will still need to obtain an account on the terminal server.

Requesting KA9Q Software

There are currently two OIRM sources offering the KA9Q package and terminal server service: the Telecommunications Services Division (TSD) in Fort Collins and the Washington Service Center (WSC) in Washington, DC. To obtain a copy of the KA9Q software package and/or an account on the terminal server, complete the form in Appendix A. Remember when choosing a service center that the serial connection is established over the public switched telephone network and may involve long distance charges. Choose that service center which will minimize your long distance charges.

Installing the KA9Q Software

After receiving the KA9Q distribution package the user must install the software. In the package there is a document describing the installation procedure as well as some basic instructions on how to use the software, a distribution diskette containing all of the software and documentation, and an informational sheet containing the terminal server userid and password. In general the installation is a very simple process involving the execution of the installation program contained on the diskette and answering some simple questions about the communications environment and the terminal server account information. The only hardware change necessary to get the KA9Q package working is to install a modem and connect it to a communications port of the PC. The installation procedure is described in detail in the "KA9Q User Reference Manual Supplement" included in this Guide as

Appendix F. After installing the software and verifying that the modem is installed and connected correctly, the software is ready to be used.

Using the KA9Q Software

When the KA9Q package is installed on a DOS PC, several files are created in the installation directory which contain documentation on the use of the KA9Q package. The "KA9Q User Reference Manual Supplement" (Appendix F) is one of these documents also included in the KA9Q package in hard copy form. The electronic version of the document is in the file INSTALL.DOC. It includes installation instructions as well as some basic instructions on how to establish modem connections and the use of the TCP/IP applications.

Included in the "KA9Q User Reference Manual Supplement" is a detailed description of how to use the NET.EXE program which is the heart of the KA9Q software package. The NET.EXE program provides the TCP/IP software including the common applications Telnet, FTP and SMTP. The NET.EXE program is documented in the "NET Users Reference Guide" contained in the file KA9Q.DOC.

The BM.EXE program has been included in the KA9Q distribution package to compose and read mail messages. Instructions on how to use BM.EXE are included in the "KA9Q User Reference Manual Supplement" (Appendix F). BM.EXE is also documented in greater detail in the "BM Reference guide" contained in the file BM.DOC.

If you have any problems or questions not addressed in the documentation, contact the help desks described in Section 4.

8.2 Network Services

OIRM provides several important network services. Included in these services are common services and guest host services. These services are defined in the sections that follow.

Common Services

Common services are defined as services that are offered for general use by all users connected to the USDA Internet. These services currently include access to a Domain Name Server and a USDA Internet help desk.

Domain Name Server

An important service provided to users connected to the USDA Internet is the domain name system. A key part of this system is the domain name server. Each of the USDA Internet facilities have a host configured as a domain name server as described in the following sections.

Description of the Domain Name Server. In the Internet environment, IP addresses are used to uniquely identify individual networks and hosts. An IP address is a 32 bit number written as four decimal numbers separated by decimal points or "dots." Because these numbers can be difficult to remember, the Domain Name System (DNS) was established in the Internet to map IP addresses to easier to remember, English sounding names or acronyms. A complete description of the DNS is not included in this Guide. For further reading on the subject see RFC's 1034 and 1035. In general the DNS is composed of logical divisions of the domain name space called zones. For each zone there are host machines that store and distribute information about a zone. The domain name **ag.gov** has been assigned to the zone defined as the USDA Internet.

To support the zone, **ag.gov**, two domain name servers have been installed on the USDA Internet network. The primary domain name server is located at the Telecommunications Services Division in Fort Collins. One of several secondary or backup domain name servers is located at the Washington Service Center in Washington DC. These servers will hold the domain name information for the **ag.gov** zone. As the network grows, additional sub-zones may be created and additional domain name servers added to the network.

Using the Domain Name System. To use the domain name system it is required that a host be connected to the network. The host has one of the domain name server IP addresses configured in the TCP/IP package. Most TCP/IP packages have the ability to configure a domain name server IP address. Consult your TCP/IP package documentation for details on how it implements domain name system. When deciding which domain name server to use, you should select the server that is geographically closest to your portion of the network. The domain name server at Fort Collins TSD is assigned the IP address "162.79.3.2" and the domain name server located at WSC is assigned the IP address "162.72.2.10" All domain name resolutions attempted by your host will be sent to the configured domain name server after specifying the appropriate address in your TCP/IP package configuration.

If you would like to implement a sub-zone of the zone **ag.gov**, refer to the "Request to Manage a Domain" discussed previously in Section 8.1 and contact the TSD with any questions. Remember that to implement a zone there must be a domain name server for that zone. Details required to establish a zone in the domain name system are described in RFC 1034 and RFC 1035.

Help Desk

An important service offered to the USDA Internet user community is the USDA Internet Help Desk. Currently each of the USDA Internet facilities are offering this service as described in the following sections.

Description of the Help Desk. The Internet is a relatively new resource to most of the USDA. As with anything new, users may find they have questions on accessing and using this resource. The intent of this Guide is to anticipate some questions that a new user may have. When a question arises that is not addressed in this Guide, the Help Desk is available as an additional resource.

Requesting Help Desk Service. The Help Desk is a service provided to all users who are connected to the USDA Internet. If you run into problems or have questions unanswered by this Guide, you can contact the nearest Help Desk. There are two Help Desks, one located at each of the USDA Internet facilities. The primary Help Desk is located in Fort Collins, Colorado and can be reached at 303-498-2007.

There is also a Help Desk located in Washington, DC. Its number is 202-720-1716.

If you call either Help Desk, please have the following information ready.

- 1) Your name, agency, and phone number.
- 2) Specify that it is an Internet question. (These Help Desks handle other types of problems as well as those dealing with the Internet)
- 3) Internet access problem: If it is an access problem, be prepared to explain how you get to the Internet when everything does work.

For example do you use:

- The Washington, DC USDA LAN and terminal server?
- The Fort Collins LAN and PC TCP software?
- A circuit from Chester, PA to Washington, DC?
- A modem and KA9Q software?

- 4) If you have an IP address assigned to your workstation, it would be helpful if you knew what it was. (IP addresses look something like this: 162.72.53.101.)
- 5) Host access problem: If you are having trouble reaching a public host, the Help Desk person can see if the host is responding and give you some general rules for logging in. However it may be necessary to refer your questions to your agency Internet contact. For instance, if you are having difficulty reaching your agency database in California, we can only check to see if it answers. We cannot bring it up, check login IDs, etc.

The Internet support Help Desk is a new service. There will be changes as we learn how to best meet your requirements. Feel free to let us know how we are doing and how service might be improved.

Guest Host Services

Each of the USDA network facilities have a guest host dedicated to providing services to users accessing the Internet. The current offering of these services include mailbox service, Telnet and some other TCP based applications. These services are described in greater detail in the following sections.

Mailbox

One of the guest host services offered by OIRM to support access to the USDA Internet is the establishment of a mail server for users who do not have an alternative way to send and receive Internet electronic mail. The following sections describe the mailbox service.

Description of Mailbox Service. Users who connect to the USDA Internet network serially from a PC do not have the capability of receiving mail in a reliable manner. This is because their connections are only active for limited amounts of time. The same is true for users who have workstations on a LAN who do not have access to a host configured to receive mail. To solve this problem two mail server hosts have been connected to the USDA Internet to be shared by the user community, one at the Telecommunications Services Division (TSD) in Fort Collins, Colorado and the Washington Service Center (WSC) in Washington DC. These hosts will have mail boxes for users who have no other means of sending and receiving mail. Users will need to submit requests to establish accounts and mailboxes on the hosts. Because there is a limit to the number of mailboxes that can be configured on one host, other ways to send and receive mail should be considered before requesting a host mailbox account.

Requesting Mailbox Service. To request a mailbox, complete the form in Appendix B and send it to the service center listed at the bottom of the form where you would like your mailbox to be located. When deciding which service center to use, remember that serial connections are established over telephone lines and may involve long distance charges. Choose the service center which will minimize your long distance charges.

Using Mailbox Service. To use the mail server the user will need to establish a terminal session with the mail server host. This can be accomplished using one of the two asynchronous methods described in Section 7.2 or by using the TCP/IP Telnet application directly from the network. Once a session is established with the host, sign onto the system using the userid and password supplied by one of the service centers.

The mail server is a UNIX based machine and therefore the mail capability is the UNIX mail utility. Other mail packages are being evaluated and may be offered at a later date. By entering the command "mail" at the UNIX user prompt, the mail utility is started and all of the standard mail functions are then available for use. These functions include commands that allow you to read and create mail messages. For a complete description of

the mail utility type “`man mail`” at the UNIX user prompt. Most general UNIX texts and manuals are also good resources for using the mail facility.

Telnet

One service offered by OIRM to support access to the USDA Internet is the TCP/IP application – Telnet. The following sections describe the Telnet application and the way the service is provided.

Description of Telnet. Telnet is the standard Internet terminal emulation protocol used to communicate interactively between hosts on a network. The Telnet application is described separately from the other TCP/IP based applications because it is available from two platforms at each USDA Internet facility. At each facility there is a communications server and a guest host. Each of these platforms provides a user with the Telnet application.

Users will need to use TCP/IP applications supplied by the guest host or communications server if they are connecting to the USDA Internet network without TCP/IP software on their PC through asynchronous dial-in terminal emulation connections. (See Section 7 for a discussion of terminal emulation access. Telnet is an application offered to these users.)

Requesting Telnet Service. To use the Telnet service on one of the guest hosts or communications servers located at TSD or WSC, a user must obtain an account on the platform being used. If Telnet is the only application you want to use and you are connecting using the terminal emulation asynchronous dial-in access described in Section 7.2, you will only need to request an account on a communications server. To request an account on one of the communications servers, fill out the form reproduced in Appendix A and send it to the appropriate service center.

If you will need access to other TCP/IP applications (discussed later in this Section as “Other TCP/IP Applications”) in addition to the Telnet application, you will need to request an account on a guest host. To request an account on one of the guest hosts, fill out the form reproduced in Appendix B and send it to the service center where the host you are going to use is located. When deciding which service center to use, remember that serial connections are established over telephone lines and may involve long distance charges. Choose that service center which will minimize your long distance charges.

Using Telnet. To use the Telnet application provided by the guest host or communications server, the user will need to establish a terminal session with the appropriate platform. This can be accomplished using one of the two asynchronous methods described in Section 7.2. Once a session has been established with the host or communications server, sign onto the system using the userid and password supplied by one of the service centers.

The guest host server is a UNIX based machine and therefore the Telnet capability is the UNIX system’s implementation. To use Telnet to connect to another host, enter the command “`telnet <host>`” at the UNIX user prompt where `<host>` is the

destination's IP address or domain name. For a complete description of the Telnet command type "`man telnet`" at the UNIX user prompt.

The communications server is a Cisco ASM/3-CS platform. The Telnet capability is Cisco's implementation.

Telnet Documentation. There are many implementations of the Telnet application provided by various vendors. Most of the command syntaxes and procedures for using the application are the same or similar for each implementation, but there are some differences. Because of these differences a generic tutorial on how to use the Telnet application is not practical. Each user should refer to the documentation provided by the vendor supplying the Telnet software. For users who are using the Cisco Communications Server or the UNIX guest host, documentation on each implementation of Telnet will be provided when an account on one of the systems is issued. Documentation on the KA9Q version of Telnet is included in the KA9Q software package.

Other TCP/IP Applications

In addition to Telnet and mail, other common TCP/IP based applications are provided by the guest host servers located at each of the USDA Internet facilities. These applications are described in the following sections.

Description of Other TCP/IP Applications. Users who are connecting to the network without TCP/IP software on their host through asynchronous dial-in terminal emulation connections will need to use TCP/IP applications supplied by the guest host. In addition to Telnet and mail (SMTP) there are many other useful TCP/IP based applications. The most commonly used of these applications is the FTP or File Transfer Protocol application. FTP is the Internet file transfer application used to transfer both text and binary files between hosts connected to the Internet.

Initially, FTP will be the only other TCP/IP based application offered on the guest host server. This initial offering will be enhanced as users define their requirements. Other applications currently under consideration are Finger, Whois, Archie and Gopher. As these applications are added, users will be notified. Descriptions of other TCP/IP applications may be found in Appendix G, "Learning More About the Internet," and Appendix H, "Zen and the Art of the Internet."

Requesting Other TCP/IP Applications Service. To use the FTP and other services on one of the guest hosts located at TSD or WSC, a user must obtain an account on the platform being used. To request an account on one of the guest hosts, fill out the form reproduced in Appendix B and send it to the service center where the host you are going to use is located. When deciding which service center to use, remember that dial-in terminal emulation asynchronous connections are established over telephone lines and may involve

long distance charges. Choose that service center which will minimize your long distance charges.

Using Other TCP/IP Applications. To use FTP and other applications, the user will need to establish a terminal session with the guest host server. This can be accomplished using one of the two asynchronous methods described in Section 7.2 or by using the TCP/IP Telnet application directly from the network. Once a session is established with the host, sign onto the system using the userid and password supplied by one of the service centers.

The guest host server is a UNIX based machine and so the FTP capability is UNIX's FTP. To use FTP to connect to another host enter the command "**ftp <host>**" at the UNIX user prompt where **<host>** is the destination's IP address or domain name. For a complete description of the FTP command type "**man ftp**" at the UNIX user prompt. This will display manual pages describing the use and capabilities of the FTP application.

FTP Documentation. There are many implementations of the FTP application provided by various vendors. Most of the command syntaxes and procedures for using the application are the same or similar for each implementation, but there are some differences. Because of these differences a generic tutorial on how to use the FTP application is not practical. Each user should refer to the documentation provided by the vendor supplying the FTP software. For users who are using the UNIX guest host, documentation on the implementation of FTP will be provided when an account on the system is issued. Documentation on the KA9Q version of FTP is included in the KA9Q software package.

8.3 Router Configuration Service

Another type of service available is the configuration of routers. This is a service offered to users who would like assistance in configuring the software running on a router used to connect to the USDA Internet.

Description of Router Configuration Service

Router configuration services will be provided by the Telecommunications Services Division staff and the Washington Service Center staff if required. Initial installation of router software and ongoing maintenance of existing router software and configurations would be some of the services provided. Cost recovery agreements may be necessary depending upon the level of service provided.

Requesting Router Configuration Service

To request the router configuration service the Agency Internet coordinator should contact the Telecommunications Services Division office in Fort Collins, Colorado or the Washington Service Center in Washington, DC.

8.4 Engineering Services

In addition to administration, network, and router services, OIRM/TSD is offering Engineering Services. Engineering services offer agencies assistance in the design and implementation of a system to be used for access to the USDA Internet.

Description of Engineering Services

Engineering services will be provided by the Fort Collins Telecommunications Services Division staff on request. Selection of routers, router management, routing protocols, addressing, size and type of circuits and FTS2000 T1 channelization would be some of the services provided. Cost recovery agreements may be necessary depending upon the level of service provided.

Requesting Engineering Services

The Agency Internet coordinator should contact:

Chief, Network Management and Analysis Branch
Telecommunications Services Division
OIRM/NCC
3825 E. Mulberry St.
Fort Collins, CO 80524
303/498-1633 FAX 303/498-1660

9. Cost Recovery

For fiscal year 1993 the cost of operating the Internet services described in this document will be absorbed by OIRM. In fiscal year 1994 a method of recovering operating costs based on usage of the network and associated services will be developed and implemented. Because the cost of this aggregated Internet access will be divided among many units, the costs will be minimized for individual units.

10. USDA Resource Center Usage

This Section describes some of the services that are available on the USDA Internet.

10.1 NCC Applications

The National Computer Center (NCC) is currently in the process of implementing TCP/IP applications on the IBM MVS Host Complex. Applications currently targeted for USDA-wide availability include: Telnet (TN3270), FTP, NFS and SMTP. These applications are currently being tested and customized and are expected to be available for production by late 1993. The NCC is also in the process of testing applications developed for CICS and Oracle. Other TCP/IP applications may be considered in the future.

For current information concerning the NCC's TCP/IP Implementation and the availability of TCP/IP applications, refer to the NCC Customer Handbook or contact the NCC Customer Support Center in Kansas City at (816) 926-6681.

10.2 Other USDA Resource Centers

As other USDA Resource Centers become available for use by the USDA Internet user community, they will be described in this section. Examples of USDA resource centers include:

- The SCS SOILS database at Ames, IA
- The SCS PLANTS database at the NCC and TASC

APPENDIX A

USDA Internet Asynchronous Access Request Form

This form is to be used to request a copy of the KA9Q SLIP software package and/or an account on the terminal server located at one of the service centers offering the USDA terminal server service. Please complete the form and send it to the appropriate service center.

Agency and Organization Name:

Contact person:

E-Mail Address:

Telephone number:

Address:

Do you want a domain name for your account and IP address?:

Yes No

If Yes, specify what it should be (One will be assigned to you if you do not have a preference. Contact your agency to see if a domain name has been assigned for your agency):

.ag.gov

Do you want to receive the KA9Q software package? Yes No

If Yes, select the type of media to use for the distribution:

- 5 1/4" Double Sided, Double Density Diskette
- 5 1/4" Double Sided, High Density Diskette
- 3 1/2" Double Sided, High Density Diskette
- Other: _____

Send completed forms to the appropriate service center:

E-Mail request to "A76ADDR@ATTMAIL.COM" or "addman@ag.gov"
or mail or FAX to:

OIRM/TSD
3825 East Mulberry St.
Fort Collins, CO 80524
FAX-(303)498-1660

Washington Service Center
14th and Independence Ave. SW
Room South -162
Washington, DC 20250
FAX 202/720-8274

APPENDIX B

USDA Internet Guest Host Services Request Form

This form is to be used to request an account on the guest host located at one of the service centers offering the USDA guest host services. These services are mailbox, TELNET and limited FTP. Please complete the form and send it to the appropriate service center.

Agency and Organization Name:

Contact person:

E-Mail Address:

Telephone number:

Address:

Describe why you need an account on the guest host, i.e., this is the only method available to you to access the USDA Internet.

Send completed forms to the appropriate service center:

E-Mail request to "A76ADDR@ATTMAIL.COM" or "addman@ag.gov"
or mail or FAX to:

OIRM/TSD
3825 East Mulberry St.
Fort Collins, CO 80524
FAX-(303)498-1660

Washington Service Center
14th and Independence Ave. SW
Room South -162
Washington, DC 20250
FAX 202/720-8274

APPENDIX C

IP Subnetwork Address Request Form

Agency Name:

Unit Name:

Location of subnetwork:

Administrative Authority:

Telephone number:

E-Mail Address:

Postal Address:

Technical Contact:

Telephone number:

E-Mail Address:

Postal Address:

Description of subnetwork-

Router on subnetwork (y/n):

If yes, Manufacturer:

Model:

Routing Protocol(s) supported-

(RIP, OSPF, IS-IS, etc.):

Type of subnetwork (Ethernet, Token Ring, etc.):

Number of hosts or devices connected to subnetwork-

Today: Within one year:

Unix Hosts: / DOS Hosts: /

Other Hosts/devices (specify): /

Is subnetwork connected to the Internet? (y/n):

If yes, describe connection point and method:

E-Mail request to "A76ADDR@ATTMAIL.COM" or "addman@ag.gov"

Provide any additional, available documentation, diagrams, and/or pictures which further describe your subnetwork(s) to:

OIRM/TSD
3825 East Mulberry St.
Fort Collins, CO 80524
FAX-(303)498-1660

APPENDIX D

USDA Domain Name Request Form

The requested name will be entered in the USDA DNS servers and will be associated with the given IP address within three working days of the receipt of the request. Confirmation of this action will be provided via E-Mail to the administrative authority listed below.

Agency Name:
Unit Name:

Administrative Authority:
E-mail Address:
Telephone number:
Address:

Technical Contact:
E-mail Address:
Telephone number:
Address:

Requested Name:

IP Address:

E-Mail request to "A76ADDR@ATTMAIL.COM" or "addman@ag.gov"
or mail or FAX to:

OIRM/TSD
3825 East Mulberry St.
Fort Collins, CO 80524
FAX-(303)498-1660

APPENDIX E

USDA Domain Management Request Form

Administrative Authority:

 E-Mail Address:

 Telephone number:

 Address:

Technical Contact:

 E-Mail Address:

 Telephone number:

 Address:

Requested Domain Name:

Expected Operational Date:

Agency Name:

 Unit Name:

USDA domains must be supported by two independent DNS servers that are located in different physical locations. The following information must be provided for these two servers.

Primary DNS Server:

 Hostname:

 IP Address:

 Hardware:

 Software:

Secondary DNS Server:

 Hostname:

 IP Address:

 Hardware:

 Software:

List any currently registered domain names that are supported under the second level domain of "ag.gov" that will be supported in the requested domain in the future. These host entries will be removed from the "ag.gov" domain name server tables on the expected operational date.

Domain names / IP addresses:

E-Mail request to "A76ADDR@ATTMAIL.COM" or "addman@ag.gov" or mail or FAX to:

OIRM/TSD
3825 East Mulberry St.
Fort Collins, CO 80524
FAX-(303)498-1660

APPENDIX F

KA9Q User Reference Manual Supplement

KA9Q User Reference Manual Supplement

1. INTRODUCTION

The MS-DOS version of the KA9Q Internet package is the result of several years of development by Phil Karn (KA9Q is his amateur radio operator identification). It was written to allow PC access to the Internet. By executing the NET.EXE program a user has access to FTP, TELNET, MAIL and most of the other common TCP/IP applications.

This document is intended to be a supplement to the *NET User Reference Manual* written by Phil Karn (KA9Q). The reference manual can be found on the distribution disk included with this package in the file KA9Q.DOC. The information included in this document is to assist the user in getting started using the MS-DOS executable file NET.EXE. Topics not addressed here may be found in the reference manual. The file KA9QIDX.DOC contains an index for the reference manual.

2. INSTALLATION

The KA9Q software package may be installed in any directory on your PC. The directory C:\KA9Q is the default for the installation program and will be used in all examples and instructions in this document. If you want to use a different directory, substitute all occurrences of C:\KA9Q with the new directory.

The only hardware change necessary to get the KA9Q package working is to connect your modem to a communications port (COM1 is the default) of the PC (The modem should be connected to a phone line).

In your installation package you should have found this document, the USDA KA9Q Distribution Information Sheet and a floppy disk (two disks if you are using low density disks). Verify that the serial number printed on the disk matches the serial number printed on the information sheet. If the numbers do not match, call the distribution center before installing the software.

To install the KA9Q software run the INSTALL program located on the installation disk included with the distribution package (If you received low density disks use disk number one). Insert the disk into your floppy drive (drive A is used in this example) and type the following command:

```
A:\install
```

The install program will prompt you to enter the drive you will be using for the installation source drive. Enter a character corresponding to the name of the drive you are using and press enter (Drive A will be used if you only press <enter>).

```
ENTER SOURCE DRIVE (DEFAULT IS A): <enter>
```

Next you will be asked to supply the name of a directory where you want the software to be installed. Enter the full path name of the directory and press <enter> (C:\KA9Q will be used if you only press <enter>). If the directory does not exist it will be created for you.

ENTER DESTINATION DIRECTORY (DEFAULT C:\KA9Q) : <enter>

Next you will be asked to supply your terminal server userid and password printed on the information sheet contained in the installation package. The password you enter will not be echoed so you will be asked to enter it twice to verify that it is correct.

ENTER USER ID SUPPLIED IN INSTALLATION PACKET: userid <enter>
ENTER PASSWORD FOR USER ID: ***** <enter>
RE-ENTER PASSWORD: ***** <enter>

Next you will be asked to enter the number of the communications port you are using on your PC. For example enter a '1' if you are using COM1 or a '2' if you are using COM2 (COM1 will be assumed if you only press <enter>).

ENTER THE NUMBER OF THE COMM PORT YOU ARE USING (DEFAULT IS 1):
<enter>

Next you will be asked to enter the speed your modem will run at. Possible values are 1200, 2400, 4800 and 9600 bps (9600 will be assumed if you only press <enter>).

ENTER THE SPEED OF YOUR MODEM (DEFAULT IS 9600) : <enter>

Next you will be asked to enter the access number(s) you must dial to connect to the telephone network you are using. This number will be inserted into the modem dialer string before the number of the terminal server you are using. For example if the telephone number of the terminal server is 111-222-3333 and you must dial 8 to get access to a telephone network outside your office the dialer string is "atdt 8,1112223333". If you need to modify this string after this package has been installed, edit the file "MODEM" created in the installation directory. In the example the number 8 must be dialed to get access to the FTS telephone network.

ENTER THE ACCESS NUMBER(S) YOU MUST DIAL TO CONNECT TO THE TELEPHONE NETWORK YOU ARE USING (PRESS <ENTER> FOR NONE) : 8 <enter>

Next you will be asked to verify that the installation disk is in the drive you have specified. Verify that the disk is in the correct drive and press <enter>.

VERIFY THAT THE INSTALLATION DISK IS IN DRIVE A: AND PRESS ENTER.
<enter>

At this point the installation program will install the KA9Q software creating the necessary directories and files for the installation. If you are using low density disks you will be prompted to insert a second disk. Once the program has completed you may use the KA9Q software.

3. USAGE

NET.EXE is the file that contains the KA9Q program. To Execute the NET.EXE program change directories to C:\KA9Q (You may want to add C:\KA9Q to your search path) and type the following command:

```
c:\ka9q> ka9q
```

Once the program is loaded enter the following NET command to dial the modem and establish the session with the Internet gateway:

```
NET> dialer s10 modem
```

Where "modem" is the name of the modem command file. The modem command file should work for most modems but if it doesn't work you may need to edit the file and make some adjustments. The dialer command is described in section 3.15 and section 6 of the reference manual.

Once the modem is successfully dialed into the Internet gateway and the SLIP session is established, you may enter any of the available NET commands (FTP, TELNET, etc). To get a complete list of the available commands type "help" or "?". All of the commands are documented in section 3 of the reference manual.

When you are finished using the Internet gateway you will need to hang up the modem to disconnect from the Internet gateway. To hang up the modem execute the following command:

```
NET> dialer s10 hangup
```

And to exit the NET program type the following command:

```
NET> exit
```

4. MAIL

Because KA9Q is designed to work on a temporary async connection the mail application is used mainly to send mail messages. The KA9Q NET.EXE program can only send the message that has been previously created so a separate program must be used to generate the message. The program supplied with this distribution to generate mail messages is the BM.EXE program.

The BM.EXE mail user interface program was created by Bdale Garbee and subsequent versions were written by Gerard van der Grinten and Dave Trulli. BM ('BM' stands for Bdale's mailer) provides a full set of mail services to the user which allow sending and receiving electronic mail. To run the BM program cd to the C:\KA9Q directory and type the following command:

c:\KA9Q> mail

This will start the BM.EXE program which you may use to create a mail message. A list of the available commands may be generated by typing a "?". The BM.EXE program and it's commands are documented in the *BM User Manual* located in the BM.DOC file provided in the installation.

Once you have created one or more mail messages you will need to use the NET.EXE program to send it. To do this run the KA9Q NET.EXE program and connect to the terminal server as described in the USAGE section of this document. Once the connection is established type the following command to send the mail messages:

NET> smtp kick

Other smtp commands are documented in section 3.55 of the KA9Q documentation.

5. FILES

Files created during the installation:

AUTOEXEC.NET	-	This is the default start-up file for the NET program. All lines in this file are executed as NET commands at the time the NET program is executed. You may specify a different start-up file on the command line of the NET program (See section 1.2.4 of the reference manual).
MODEM	-	This file contains modem commands to dial the modem and commands to sign onto the Internet gateway system. This file may need to be adjusted for your own modem and system.
HANGUP	-	This file contains modem command to hangup the modem.
NET.EXE	-	The KA9Q executable program.
KA9Q.BAT	-	Command file to execute the NET.EXE program.
KA9Q.DOC	-	The NET User Reference Manual.
KA9QIDX.DOC	-	The index for the NET User Reference Manual.
INSTALL.DOC	-	Installation instructions and User Manual Supplement.
BM.EXE	-	The BM executable program.
MAIL.BAT	-	Command file to execute the BM.EXE program.

- BM.DOC - The BM User Manual.
- BM.RC - The BM configuration file (located in the root directory).

APPENDIX G

Learning More About the Internet

Learning More About the Internet

This Appendix is written for the new or inexperienced user of the Internet. Included in the following text are instructions on how to access an online database of information on all kinds of topics relating to the Internet. A prerequisite requirement to accessing this information is access to the Internet and more specifically, access to the Telnet application on the Internet. If you do not have this access, contact your network administrator for details on how to get access to the Internet from your network and access to the Telnet application. If you do not have a network administrator or someone acting in that role, contact one of the help desks described in section 8.2 for assistance.

Once you have access to the Internet and a version of the Telnet application, you may access one of many information servers running an application called Gopher. Gopher is a menu driven distributed information delivery system used in the Internet to deliver information on various topics to the Internet user community. To access a Gopher server, a Gopher client is used. The client we are describing is located on a host managed by the InterNIC, and is available to anyone connected to the Internet. There are many other Gopher sites around the Internet that offer a wide variety of information. In the future a Gopher server will be established for the USDA Internet with information particular to that network.

To access the InterNIC's gopher server, Telnet to rs.internic.net and type gopher. Starting from the system you are using to connect to the Internet, the Telnet command sequence should look something like this:

```
> telnet rs.internic.net
```

If you get an error message stating that the domain name rs.internic.net does not exist, try the following command:

```
> telnet 198.41.0.5
```

If this fails, contact your network administrator for assistance.

After a successful Telnet the following screen will be displayed:

```
ag.gov> telnet rs.internic.net
Trying 198.41.0.5 ...
Connected to rs.internic.net.
Escape character is '^]'.

SunOS UNIX (rs) (ttyq9)

*****
* -- InterNIC Registration Services Center --
*
* For gopher, type:          GOPHER <return>
* For wais, type:            WAIS <search string> <return>
* For the *original* whois type: WHOIS [search string] <return>
* For the X.500 whois DUA, type: X500WHOIS <return>
* For registration status:   STATUS <ticket number> <return>
*
* For user assistance call (800) 444-4345 | (619) 455-4600 or (703)
* 742-4777
* Please report system problems to ACTION@rs.internic.net
*****
Please be advised that the InterNIC Registration host contains INTERNET
Domains, IP Network Numbers, ASNs, and Points of Contacts ONLY. Please
refer to rfc1400.txt for details (available via anonymous ftp at
either nic.ddn.mil [/rfc/rfc1400.txt] or ftp.rs.internic.net
[/policy/rfc1400.txt]).
```

Cmdinter Ver 1.3 Fri Aug 20 14:52:11 1993 EST
[term-type] InterNIC >

At the InterNIC prompt type "gopher" to start the Gopher client:

[term-type] InterNIC > **gopher**

The following screen will then be displayed:

Internet Gopher Information Client v1.03


Root gopher server: rs.internic.net

- > 1. Information about the InterNIC/
- 2. InterNIC Information Services (General Atomics)/
- 3. InterNIC Registration Services (NSI)/
- 4. InterNIC Directory and Database Services (AT&T)/

Press ? for Help, q to Quit, u to go up a menu

Page: 1/1

At this point you are free to chose any topic you would like to get information on. For the first time user a good place to start is option 2, "InterNIC Information Services (General Atomics)". To select this option use the down arrow key to move the "-->" to option 2 and press <return>. After selecting this option the following screen will be displayed:



From here select the topic you are interested in and continue to make menu selections until the document you are interested in is displayed. If while viewing a document you wish to stop, press "q" and follow the directions at the bottom of the screen. While navigating the menu system, command options are also listed at the bottom of the screen.

To exit the Gopher client type "q" and verify that you want to exit the InterNIC Gopher client. Next the InterNIC host prompt will be displayed. At this prompt enter "exit" to logoff the host:

InterNIC > **exit**

APPENDIX H

Zen and the Art of the Internet

Editor's Note. This paper by Brendan P. Kehoe is reproduced by kind permission of the author.

Zen and the Art of the Internet

A Beginner's Guide to the Internet
First Edition
January 1992

by **Brendan P. Kehoe**

This is revision 1.0 of February 2, 1992.

Copyright © 1992 Brendan P. Kehoe

Permission is granted to make and distribute verbatim copies of this guide provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this booklet under the conditions for verbatim copying, provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this booklet into another language, under the above conditions for modified versions, except that this permission notice may be stated in a translation approved by the author.

Table of Contents

Preface	1
Acknowledgements	3
1 Network Basics	5
1.1 Domains	5
1.2 Internet Numbers.....	7
1.3 Resolving Names and Numbers	7
1.4 The Networks	8
1.5 The Physical Connection	8
2 Electronic Mail	11
2.1 Email Addresses.....	11
2.1.1 %@!.. Symbolic Cacophony	11
2.1.2 Sending and Receiving Mail	12
2.1.3 Anatomy of a Mail Header	13
2.1.4 Bounced Mail	14
2.2 Mailing Lists	15
2.2.1 Listservs	16
3 Anonymous FTP	19
3.1 FTP Etiquette	19
3.2 Basic Commands	20
3.2.1 Creating the Connection.....	20
3.2.2 dir	21
3.2.3 cd	22
3.2.4 get and put	22
3.2.4.1 ASCII vs Binary	23
3.2.4.2 mget and mput	24
3.3 The archie Server.....	25
3.3.1 Using archie Today.....	25
3.3.2 archie Clients	26
3.3.3 Mailing archie.....	27
3.3.4 The whatis database	27
4 Usenet News.....	29
4.1 What Usenet Is	29

4.2	The Diversity of Usenet	29
4.3	What Usenet Is Not	29
4.4	Propagation of News	31
4.5	Group Creation	32
4.6	If You're Unhappy.....	33
4.7	The History of Usenet (The ABCs)	33
4.8	Hierarchies	34
4.9	Moderated vs Unmoderated	35
4.10	<i>news.groups</i> & <i>news.announce.newgroups</i>	36
4.11	How Usenet Works	36
4.12	Mail Gateways	37
4.13	Usenet "Netiquette"	37
4.13.1	Signatures	37
4.13.2	Posting Personal Messages.....	38
4.13.3	Posting Mail	38
4.13.4	Test Messages	39
4.13.5	Famous People Appearing.....	39
4.13.6	Summaries	39
4.13.7	Quoting.....	40
4.13.8	Crossposting	41
4.13.9	Recent News	41
4.13.10	Quality of Postings.....	41
4.13.11	Useful Subjects	42
4.13.12	Tone of Voice	42
4.13.13	Computer Religion	43
4.14	Frequently Asked Questions	43
4.14.1	The Pit-Manager Archive	43
5	Telnet	45
5.1	Using Telnet	45
5.1.1	Telnet Ports.....	45
5.2	Publicly Accessible Libraries.....	46
5.3	The Cleveland Freenet.....	47
5.4	Directories	47
5.4.1	Knowbot	48
5.4.2	White Pages	48
5.5	Databases	48
5.5.1	Colorado Alliance of Research Libraries (CARL) ..	48
5.5.2	PENpages.....	49
5.5.3	Clemson Univ. Forestry & Agricultural Network ..	49
5.5.4	University of Maryland Info Database	49
5.5.5	University of Michigan Weather Underground	50
5.5.6	Geographic Name Server	50
5.5.7	FEDIX—Minority Scholarship Information	50

5.5.8	Science & Technology Information System	51
5.5.9	Ocean Network Information Center.....	51
5.5.10	NASA/IPAC Extragalactic Database (NED).....	51
5.5.11	U.S. Naval Observatory Automated Data Service	52
6	Various Tools	53
6.1	Finger.....	53
6.2	Ping	54
6.3	Talk.....	55
6.4	The WHOIS Database.....	55
6.4.1	Other Uses of WHOIS.....	57
7	Commercial Services	59
7.1	Electronic Journals	59
7.2	Commercial Databases	60
7.3	Clarinet News	60
8	Things You'll Hear About	63
8.1	The Internet Worm.....	63
8.2	The Cuckoo's Egg	64
8.3	Organizations	64
8.3.1	The Association for Computing Machinery.....	65
8.3.2	Computer Professionals for Social Responsibility ..	65
8.3.3	The Electronic Frontier Foundation	66
8.3.4	The Free Software Foundation	68
8.3.5	The League for Programming Freedom.....	68
8.4	Networking Initiatives	69
8.4.1	NREN	69
9	Finding Out More.....	71
9.1	Internet Resource Guide.....	71
9.2	Requests for Comments.....	71
Conclusion	73	
Appendix A	Getting to Other Networks	75
Appendix B	Retrieving Files via Email.....	77
Archive Servers.....	77	
FTP-by-Mail Servers	77	

Appendix C Newsgroup Creation	79
Discussion	79
Voting	79
The Result of a Vote	81
Creation of the Group	81
Glossary	83
Bibliography	91
Books	91
Periodicals & Papers	92
Index	95

Preface

The composition of this booklet was originally started because the Computer Science department at Widener University was in desperate need of documentation describing the capabilities of this “great new Internet link” we obtained.

It’s since grown into an effort to acquaint the reader with much of what’s currently available over the Internet. Aimed at the novice user, it attempts to remain operating system “neutral”—little information herein is specific to Unix, VMS, or any other environment. This booklet will, hopefully, be usable by nearly anyone.

Some typographical conventions are maintained throughout this guide. All abstract items like possible filenames, usernames, etc., are all represented in *italics*. Likewise, definite filenames and email addresses are represented in a quoted ‘typewriter’ font. A user’s session is usually offset from the rest of the paragraph, as such:

`prompt> command`

The results are usually displayed here.

The purpose of this booklet is two-fold: first, it’s intended to serve as a reference piece, which someone can easily grab on the fly and look something up. Also, it forms a foundation from which people can explore the vast expanse of the Internet. *Zen and the Art of the Internet* doesn’t spend a significant amount of time on any one point; rather, it provides enough for people to learn the specifics of what his or her local system offers.

One warning is perhaps in order—this territory we are entering can become a fantastic time-sink. Hours can slip by, people can come and go, and you’ll be locked into Cyberspace. Remember to do your work!

With that, I welcome you, the new user, to The Net.

brendan@cs.widener.edu
Chester, PA

Acknowledgements

Certain sections in this booklet are not my original work—rather, they are derived from documents that were available on the Internet and already aptly stated their areas of concentration. The chapter on Usenet is, in large part, made up of what's posted monthly to `news.announce.newusers`, with some editing and rewriting. Also, the main section on `archie` was derived from '`what is .archie`' by Peter Deutsch of the McGill University Computing Centre. It's available via anonymous FTP from `archie.mcgill.ca`. Much of what's in the telnet section came from an impressive introductory document put together by SuraNet. Some definitions in the one are from an excellent glossary put together by Colorado State University.

This guide would not be the same without the aid of many people on The Net, and the providers of resources that are already out there. I'd like to thank the folks who gave this a read-through and returned some excellent comments, suggestions, and criticisms, and those who provided much-needed information on the fly. Glee Willis deserves particular mention for all of his work; this guide would have been considerably less polished without his help.

- Andy Blankenbiller, Army at Aberdeen
- Alan Emtage, McGill University Computer Science Department
- Brian Fitzgerald, Rensselaer Polytechnic Institute
- John Goetsch, Rhodes University, South Africa
- Jeff Kellem, Boston University's Chemistry Department
- Bill Krauss, Moravian College
- Steve Lordin, Delco Electronics
- Mike Nesel, NASA
- Bob Neveln, Widener University Computer Science Department
- Wanda Pierce, McGill University Computing Centre
- Joshua Poulson, Widener University Computing Services
- Dave Sill, Oak Ridge National Laboratory
- Bob Smart, CitiCorp/TTI
- Ed Vielmetti, Vice President of MSEN
- Craig Ward, USC/Information Sciences Institute (ISI)
- Glee Willis, University of Nevada, Reno
- Chip Yamasaki, OSHA

1 Network Basics

We are truly in an information society. Now more than ever, moving vast amounts of information quickly across great distances is one of our most pressing needs. From small one-person entrepreneurial efforts, to the largest of corporations, more and more professional people are discovering that the only way to be successful in the '90s and beyond is to realize that technology is advancing at a break-neck pace—and they must somehow keep up. Likewise, researchers from all corners of the earth are finding that their work thrives in a networked environment. Immediate access to the work of colleagues and a "virtual" library of millions of volumes and thousands of papers affords them the ability to incorporate a body of knowledge heretofore unthinkable. Work groups can now conduct interactive conferences with each other, paying no heed to physical location—the possibilities are endless.

You have at your fingertips the ability to talk in "real-time" with someone in Japan, send a 2,000-word short story to a group of people who will critique it for the sheer pleasure of doing so, see if a Macintosh sitting in a lab in Canada is turned on, and find out if someone happens to be sitting in front of their computer (*logged on*) in Australia, all inside of thirty minutes. No airline (or tardis, for that matter) could ever match that travel itinerary.

The largest problem people face when first using a network is grasping all that's available. Even seasoned users find themselves surprised when they discover a new service or feature that they'd never known even existed. Once acquainted with the terminology and sufficiently comfortable with making occasional mistakes, the learning process will drastically speed up.

1.1 Domains

Getting where you want to go can often be one of the more difficult aspects of using networks. The variety of ways that places are named will probably leave a blank stare on your face at first. Don't fret; there is a method to this apparent madness.

If someone were to ask for a home address, they would probably expect a street, apartment, city, state, and zip code. That's all the information the post office needs to deliver mail in a reasonably speedy fashion. Likewise, computer addresses have a structure to them. The general form is:

a person's email address on a computer: `user@somewhere.domain`
a computer's name: `somewhere.domain`

The user portion is usually the person's account name on the system, though it doesn't have to be. `somewhere.domain` tells you the name of a

system or location, and what kind of organization it is. The trailing *domain* is often one of the following:

com	Usually a company or other commercial institution or organization, like Convex Computers ('convex.com').
edu	An educational institution, e.g. New York University, named 'nyu.edu'.
gov	A government site; for example, NASA is 'nasa.gov'.
mil	A military site, like the Air Force ('af.mil').
net	Gateways and other administrative hosts for a network (it does not mean all of the hosts in a network). ¹ One such gateway is 'near.net'.
org	This is a domain reserved for private organizations, who don't comfortably fit in the other classes of domains. One example is the Electronic Frontier Foundation (see Section 8.3.3 [EFF], page 66), named 'eff.org'.

Each country also has its own top-level domain. For example, the *us* domain includes each of the fifty states. Other countries represented with domains include:

au	Australia
ca	Canada
fr	France
uk	The United Kingdom. These also have sub-domains of things like 'ac.uk' for academic sites and 'co.uk' for commercial ones.

The proper terminology for a site's domain name (*somewhere.domain* above) is its *Fully Qualified Domain Name* (FQDN). It is usually selected to give a clear indication of the site's organization or sponsoring agent. For example, the Massachusetts Institute of Technology's FQDN is 'mit.edu'; similarly, Apple Computer's domain name is 'apple.com'. While such obvious names are usually the norm, there are the occasional exceptions that are ambiguous enough to mislead—like 'vt.edu', which on first impulse one might surmise is an educational institution of some sort in Vermont; not so. It's actually the domain name for Virginia Tech. In most cases it's relatively easy to glean the meaning of a domain name—such confusion is far from the norm.

¹ *The Matrix*, 111.

1.2 Internet Numbers

Every single machine on the Internet has a unique address,² called its *Internet number* or *IP Address*. It's actually a 32-bit number, but is most commonly represented as four numbers joined by periods ('.'), like 147.31.254.130. This is sometimes also called a *dotted quad*; there are literally thousands of different possible dotted quads. The ARPAnet (the mother to today's Internet) originally only had the capacity to have up to 256 systems on it because of the way each system was addressed. In the early eighties, it became clear that things would fast outgrow such a small limit; the 32-bit addressing method was born, freeing thousands of host numbers.

Each piece of an Internet address (like 192) is called an "octet," representing one of four sets of eight bits. The first two or three pieces (e.g. 192.55.239) represent the network that a system is on, called its *subnet*. For example, all of the computers for Wesleyan University are in the subnet 129.133. They can have numbers like 129.133.10.10, 129.133.230.19, up to 65 thousand possible combinations (possible computers).

IP addresses and domain names aren't assigned arbitrarily—that would lead to unbelievable confusion. An application must be filed with the Network Information Center (NIC), either electronically (to `hostmaster@nic.ddn.mil`) or via regular mail.

1.3 Resolving Names and Numbers

Ok, computers can be referred to by either their FQDN or their Internet address. How can one user be expected to remember them all?

They aren't. The Internet is designed so that one can use either method. Since humans find it much more natural to deal with words than numbers in most cases, the FQDN for each host is mapped to its Internet number. Each domain is served by a computer within that domain, which provides all of the necessary information to go from a domain name to an IP address, and vice-versa. For example, when someone refers to `foosun.bar.com`, the resolver knows that it should ask the system `foovax.bar.com` about systems in `bar.com`. It asks what Internet address `foosun.bar.com` has; if the name `foosun.bar.com` really exists, `foovax` will send back its number. All of this "magic" happens behind the scenes.

² At least one address, possibly two or even three—but we won't go into that.

Rarely will a user have to remember the Internet number of a site (although often you'll catch yourself remembering an apparently obscure number, simply because you've accessed the system frequently). However, you will remember a substantial number of FQDNs. It will eventually reach a point when you are able to make a reasonably accurate guess at what domain name a certain college, university, or company might have, given just their name.

1.4 The Networks

Internet The Internet is a large "network of networks." There is no one network known as The Internet; rather, regional nets like SuraNet, PrepNet, NearNet, et al., are all inter-connected (nay, "inter-networked") together into one great living thing, communicating at amazing speeds with the TCP/IP protocol. All activity takes place in "real-time."

UUCP The UUCP network is a loose association of systems all communicating with the 'UUCP' protocol. (UUCP stands for 'Unix-to-Unix Copy Program'.) It's based on two systems connecting to each other at specified intervals, called *polling*, and executing any work scheduled for either of them. Historically most UUCP was done with Unix equipment, although the software's since been implemented on other platforms (e.g. VMS). For example, the system *oregano* polls the system *basil* once every two hours. If there's any mail waiting for *oregano*, *basil* will send it at that time; likewise, *oregano* will at that time send any jobs waiting for *basil*.

BITNET BITNET (the "Because It's Time Network") is comprised of systems connected by *point-to-point* links, all running the NJE protocol. It's continued to grow, but has found itself suffering at the hands of the falling costs of Internet connections. Also, a number of mail gateways are in place to reach users on other networks.

1.5 The Physical Connection

The actual connections between the various networks take a variety of forms. The most prevalent for Internet links are *56k leased lines* (dedicated telephone lines carrying 56kilobit-per-second connections) and *T1 links* (special phone lines with 1Mbps connections). Also installed are *T3 links*, acting

as backbones between major locations to carry a massive 45Mbps load of traffic.

These links are paid for by each institution to a local carrier (for example, Bell Atlantic owns PrepNet, the main provider in Pennsylvania). Also available are *SLIP* connections, which carry Internet traffic (packets) over high-speed modems.

UUCP links are made with modems (for the most part), that run from 1200 baud all the way up to as high as 38.4Kbps. As was mentioned in Section 1.4 [The Networks], page 8, the connections are of the *store-and-forward* variety. Also in use are Internet-based UUCP links (as if things weren't already confusing enough!). The systems do their UUCP traffic over TCP/IP connections, which give the UUCP-based network some blindingly fast "hops," resulting in better connectivity for the network as a whole. UUCP connections first became popular in the 1970's, and have remained in wide-spread use ever since. Only with UUCP can Joe Smith correspond with someone across the country or around the world, for the price of a local telephone call.

BITNET links mostly take the form of 9600bps modems connected from site to site. Often places have three or more links going; the majority, however, look to "upstream" sites for their sole link to the network.

"The Glory and the Nothing of a Name"
Byron, *Churchill's Grave*

2 Electronic Mail

The desire to communicate is the essence of networking. People have always wanted to correspond with each other in the fastest way possible, short of normal conversation. Electronic mail (or *email*) is the most prevalent application of this in computer networking. It allows people to write back and forth without having to spend much time worrying about how the message actually gets delivered. As technology grows closer and closer to being a common part of daily life, the need to understand the many ways it can be utilized and how it works, at least to some level, is vital.

2.1 Email Addresses

Electronic mail is hinged around the concept of an address; the section on Networking Basics made some reference to it while introducing domains. Your *email address* provides all of the information required to get a message to you from anywhere in the world. An address doesn't necessarily have to go to a human being. It could be an archive server,¹ a list of people, or even someone's pocket pager. These cases are the exception to the norm—mail to most addresses is read by human beings.

2.1.1 %@!..: Symbolic Cacophony

Email addresses usually appear in one of two forms—using the Internet format which contains '@', an "at"-sign, or using the UUCP format which contains '!', an exclamation point, also called a "bang." The latter of the two, UUCP "bang" paths, is more restrictive, yet more clearly dictates how the mail will travel.

To reach Jim Morrison on the system `south.america.org`, one would address the mail as '`jm@south.america.org`'. But if Jim's account was on a UUCP site named `brazil`, then his address would be '`brazil!jm`'. If it's possible (and one exists), try to use the Internet form of an address; bang paths can fail if an intermediate site in the path happens to be down. There is a growing trend for UUCP sites to register Internet domain names, to help alleviate the problem of path failures.

Another symbol that enters the fray is '%'—it acts as an extra "routing" method. For example, if the UUCP site `dream` is connected to

¹ See [Archive Servers], page 77, for a description.

south.america.org, but doesn't have an Internet domain name of its own, a user debbie on dream can be reached by writing to the address

`debbie%dream@south.america.org`

The form is significant. This address says that the local system should first send the mail to south.america.org. There the address `debbie%dream` will turn into `debbie@dream`, which will hopefully be a valid address. Then south.america.org will handle getting the mail to the host dream, where it will be delivered locally to debbie.

All of the intricacies of email addressing methods are fully covered in the book *!%@:: A Directory of Electronic Mail Addressing and Networks* published by O'Reilly and Associates, as part of their Nutshell Handbook series. It is a must for any active email user. Write to `nuts@ora.com` for ordering information.

2.1.2 Sending and Receiving Mail

We'll make one quick diversion from being OS-neuter here, to show you what it will look like to send and receive a mail message on a Unix system. Check with your system administrator for specific instructions related to mail at your site.

A person sending the author mail would probably do something like this:

```
% mail brendan@cs.widener.edu
Subject: print job's stuck
```

I typed 'print babe.gif' and it didn't work! Why??

The next time the author checked his mail, he would see it listed in his mailbox as:

```
% mail
"/usr/spool/mail/brendan": 1 messages 1 new 1 unread
U 1 joeuser@foo.widene Tue May  5 20:36 29/956  print job's stuck
?
```

which gives information on the sender of the email, when it was sent, and the subject of the message. He would probably use the 'reply' command of Unix mail to send this response:

```
? r
To: joeuser@foo.widener.edu
Subject: Re: print job's stuck
```

You shouldn't print binary files like GIFs to a printer!

Brendan

Try sending yourself mail a few times, to get used to your system's mailer. It'll save a lot of wasted aspirin for both you and your system administrator.

2.1.3 Anatomy of a Mail Header

An electronic mail message has a specific structure to it that's common across every type of computer system.² A sample would be:

```
From bush@hq.mil Sat May 25 17:06:01 1991
Received: from hq.mil by house.gov with SMTP id AA21901
          (4.1/SMI for dan@house.gov); Sat, 25 May 91 17:05:56 -0400
Date: Sat, 25 May 91 17:05:56 -0400
From: The President <bush@hq.mil>
Message-Id: <9105252105.AA06631@hq.mil>
To: dan@senate.gov
Subject: Meeting
```

Hi Dan .. we have a meeting at 9:30 a.m. with the Joint Chiefs. Please don't oversleep this time.

The first line, with 'From' and the two lines for 'Received:' are usually not very interesting. They give the "real" address that the mail is coming from (as opposed to the address you should reply to, which may look much different), and what places the mail went through to get to you. Over the Internet, there is always at least one 'Received:' header and usually no more than four or five. When a message is sent using UUCP, one 'Received:' header is added for each system that the mail passes through. This can often result in more than a dozen 'Received:' headers. While they help with dissecting problems in mail delivery, odds are the average user will never want to see them. Most mail programs will filter out this kind of "cruft" in a header.

The 'Date:' header contains the date and time the message was sent. Likewise, the "good" address (as opposed to "real" address) is laid out in the 'From:' header. Sometimes it won't include the full name of the person

² The standard is written down in RFC-822. See [RFCs], page 73 for more info on how to get copies of the various RFCs.

(in this case 'The President'), and may look different, but it should always contain an email address of some form.

The 'Message-ID:' of a message is intended mainly for tracing mail routing, and is rarely of interest to normal users. Every 'Message-ID:' is guaranteed to be unique.

'To:' lists the email address (or addresses) of the recipients of the message. There may be a 'Cc:' header, listing additional addresses. Finally, a brief subject for the message goes in the 'Subject:' header.

The exact order of a message's headers may vary from system to system, but it will always include these fundamental headers that are vital to proper delivery.

2.1.4 Bounced Mail

When an email address is incorrect in some way (the system's name is wrong, the domain doesn't exist, whatever), the mail system will *bounce* the message back to the sender, much the same way that the Postal Service does when you send a letter to a bad street address. The message will include the reason for the bounce; a common error is addressing mail to an account name that doesn't exist. For example, writing to Lisa Simpson at Widener University's Computer Science department will fail, because she doesn't have an account.³

```
From: Mail Delivery Subsystem <MAILER-DAEMON>
Date: Sat, 25 May 91 16:45:14 -0400
To: mg@gracie.com
Cc: Postmaster@cs.widener.edu
Subject: Returned mail: User unknown
```

```
----- Transcript of session follows -----
While talking to cs.widener.edu:
>>> RCPT To:<lsimpson@cs.widener.edu>
<<< 550 <lsimpson@cs.widener.edu>... User unknown
550 lsimpson... User unknown
```

As you can see, a carbon copy of the message (the 'Cc:' header entry) was sent to the postmaster of Widener's CS department. The *Postmaster* is responsible for maintaining a reliable mail system on his system. Usually postmasters at sites will attempt to aid you in getting your mail where it's

³ Though if she asked, we'd certainly give her one.

supposed to go. If a typing error was made, then try re-sending the message. If you're sure that the address is correct, contact the postmaster of the site directly and ask him how to properly address it.

The message also includes the text of the mail, so you don't have to retype everything you wrote.

----- Unsent message follows -----

```
Received: by cs.widener.edu id AA06528; Sat, 25 May 91 16:45:14 -0400
Date: Sat, 25 May 91 16:45:14 -0400
From: Matt Groening <mg@gracie.com>
Message-Id: <9105252045.AA06528@gracie.com>
To: lsimpson@cs.widener.edu
Subject: Scripting your future episodes
Reply-To: writing-group@gracie.com
```

... verbiage ...

The full text of the message is returned intact, including any headers that were added. This can be cut out with an editor and fed right back into the mail system with a proper address, making redelivery a relatively painless process.

2.2 Mailing Lists

People that share common interests are inclined to discuss their hobby or interest at every available opportunity. One modern way to aid in this exchange of information is by using a *mailing list*—usually an email address that redistributes all mail sent to it back out to a list of addresses. For example, the Sun Managers mailing list (of interest to people that administer computers manufactured by Sun) has the address ‘sun-managers@eecs.nwu.edu’. Any mail sent to that address will “explode” out to each person named in a file maintained on a computer at Northwestern University.

Administrative tasks (sometimes referred to as *administrivia*) are often handled through other addresses, typically with the suffix ‘-request’. To continue the above, a request to be added to or deleted from the Sun Managers list should be sent to ‘sun-managers-request@eecs.nwu.edu’.

When in doubt, try to write to the ‘-request’ version of a mailing list address first; the other people on the list aren’t interested in your desire to be added or deleted, and can certainly do nothing to expedite your request. Often if the administrator of a list is busy (remember, this is all peripheral to real jobs and real work), many users find it necessary to ask again and again,

often with harsher and harsher language, to be removed from a list. This does nothing more than waste traffic and bother everyone else receiving the messages. If, after a reasonable amount of time, you still haven't succeeded to be removed from a mailing list, write to the postmaster at that site and see if they can help.

Exercise caution when replying to a message sent by a mailing list. If you wish to respond to the author only, make sure that the only address you're replying to is that person, and not the entire list. Often messages of the sort "Yes, I agree with you completely!" will appear on a list, boring the daylights out of the other readers. Likewise, if you explicitly do want to send the message to the whole list, you'll save yourself some time by checking to make sure it's indeed headed to the whole list and not a single person.

A list of the currently available mailing lists is available in at least two places; the first is in a file on `ftp.nisc.sri.com` called 'interest-groups' under the 'netinfo/' directory. It's updated fairly regularly, but is large (presently around 700K), so only get it every once in a while. The other list is maintained by Gene Spafford (`spaf@cs.purdue.edu`), and is posted in parts to the newsgroup `news.lists` semi-regularly. (See Chapter 4 [Usenet News], page 29, for info on how to read that and other newsgroups.)

2.2.1 Listservs

On BITNET there's an automated system for maintaining discussion lists called the *listserv*. Rather than have an already harried and overworked human take care of additions and removals from a list, a program performs these and other tasks by responding to a set of user-driven commands.

Areas of interest are wide and varied—ETHICS-L deals with ethics in computing, while ADND-L has to do with a role-playing game. A full list of the available BITNET lists can be obtained by writing to 'LISTSERV@BITNIC.BITNET' with a body containing the command

`list global`

However, be sparing in your use of this—see if it's already on your system somewhere. The reply is quite large.

The most fundamental command is 'subscribe'. It will tell the listserv to add the sender to a specific list. The usage is

`subscribe foo-1 Your Real Name`

It will respond with a message either saying that you've been added to the list, or that the request has been passed on to the system on which the list is actually maintained.

The mate to 'subscribe' is, naturally, 'unsubscribe'. It will remove a given address from a BITNET list. It, along with all other listserv commands, can be abbreviated—'subscribe' as 'sub', 'unsubscribe' as 'unsub', etc. For a full list of the available listserv commands, write to 'LISTSERV@BITNIC.BITNET', giving it the command 'help'.

As an aside, there have been implementations of the listserv system for non-BITNET hosts (more specifically, Unix systems). One of the most complete is available on `cs.bu.edu` in the directory 'pub/listserv'.

"I made this letter longer than usual because
I lack the time to make it shorter."
Pascal, *Provincial Letters XVI*

3 Anonymous FTP

FTP (*File Transfer Protocol*) is the primary method of transferring files over the Internet. On many systems, it's also the name of the program that implements the protocol. Given proper permission, it's possible to copy a file from a computer in South Africa to one in Los Angeles at very fast speeds (on the order of 5-10K per second). This normally requires either a user id on both systems or a special configuration set up by the system administrator(s).

There is a good way around this restriction—the *anonymous* *FTP* service. It essentially will let anyone in the world have access to a certain area of disk space in a non-threatening way. With this, people can make files publicly available with little hassle. Some systems have dedicated entire disks or even entire computers to maintaining extensive archives of source code and information. They include *gatekeeper.dec.com* (Digital), *wuarchive.wustl.edu* (Washington University in Saint Louis), and *archive.cis.ohio-state.edu* (The Ohio State University).

The process involves the “foreign” user (someone not on the system itself) creating an *FTP* connection and logging into the system as the user ‘*anonymous*’, with an arbitrary password:

```
Name (foo.site.com:you): anonymous
Password: jm@south.america.org
```

Custom and netiquette dictate that people respond to the *Password:* query with an email address so that the sites can track the level of *FTP* usage, if they desire. (See Section 2.1 [Addresses], page 11 for information on email addresses).

The speed of the transfer depends on the speed of the underlying link. A site that has a 9600bps SLIP connection will not get the same throughput as a system with a 56k leased line (see Section 1.5 [The Physical Connection], page 8, for more on what kinds of connections can exist in a network). Also, the traffic of all other users on that link will affect performance. If there are thirty people all *FTP*ping from one site simultaneously, the load on the system (in addition to the network connection) will degrade the overall throughput of the transfer.

3.1 FTP Etiquette

Lest we forget, the Internet is there for people to do work. People using the network and the systems on it are doing so for a purpose, whether it be

research, development, whatever. Any heavy activity takes away from the overall performance of the network as a whole.

The effects of an FTP connection on a site and its link can vary; the general rule of thumb is that any extra traffic created detracts from the ability of that site's users to perform their tasks. To help be considerate of this, it's *highly* recommended that FTP sessions be held only after normal business hours for that site, preferably late at night. The possible effects of a large transfer will be less destructive at 2 a.m. than 2 p.m. Also, remember that if it's past dinner time in Maine, it's still early afternoon in California—think in terms of the current time at the site that's being visited, not of local time.

3.2 Basic Commands

While there have been many extensions to the various FTP clients out there, there is a de facto "standard" set that everyone expects to work. For more specific information, read the manual for your specific FTP program. This section will only skim the bare minimum of commands needed to operate an FTP session.

3.2.1 Creating the Connection

The actual command to use FTP will vary among operating systems; for the sake of clarity, we'll use 'FTP' here, since it's the most general form.

There are two ways to connect to a system—using its *hostname* or its Internet number. Using the hostname is usually preferred. However, some sites aren't able to resolve hostnames properly, and have no alternative. We'll assume you're able to use hostnames for simplicity's sake. The form is

```
ftp somewhere.domain
```

See Section 1.1 [Domains], page 5 for help with reading and using domain names (in the example below, *somewhere.domain* is *ftp.uu.net*).

You must first know the name of the system you want to connect to. We'll use 'ftp.uu.net' as an example. On your system, type:

```
ftp ftp.uu.net
```

(the actual syntax will vary depending on the type of system the connection's being made from). It will pause momentarily then respond with the message

```
Connected to ftp.uu.net.
```

and an initial prompt will appear:

```
220 uunet FTP server (Version 5.100 Mon Feb 11 17:13:28 EST 1991) ready.  
Name (ftp.uu.net:jm):
```

to which you should respond with anonymous:

```
220 uunet FTP server (Version 5.100 Mon Feb 11 17:13:28 EST 1991) ready.  
Name (ftp.uu.net:jm): anonymous
```

The system will then prompt you for a password; as noted previously, a good response is your email address:

```
331 Guest login ok, send ident as password.  
Password: jm@south.america.org  
230 Guest login ok, access restrictions apply.  
ftp>
```

The password itself will not echo. This is to protect a user's security when he or she is using a real account to FTP files between machines. Once you reach the `ftp>` prompt, you know you're logged in and ready to go.

3.2.2 dir

At the '`ftp>`' prompt, you can type a number of commands to perform various functions. One example is 'dir'—it will list the files in the current directory. Continuing the example from above:

```
ftp> dir  
  
200 PORT command successful.  
150 Opening ASCII mode data connection for /bin/ls.  
total 3116  
drwxr-xr-x 2 7 21 512 Nov 21 1988 .forward  
-rw-rw-r-- 1 7 11 0 Jun 23 1988 .hushlogin  
drwxrwxr-x 2 0 21 512 Jun 4 1990 Census  
drwxrwxr-x 2 0 120 512 Jan 8 09:36 ClariNet  
... etc etc ...  
-rw-rw-r-- 1 7 14 42390 May 20 02:24 newthisweek.Z  
... etc etc ...  
-rw-rw-r-- 1 7 14 2018887 May 21 01:01 uuimap.tar.Z  
drwxrwxr-x 2 7 6 1024 May 11 10:58 uunet-info  
  
226 Transfer complete.  
5414 bytes received in 1.1 seconds (4.9 Kbytes/s)  
ftp>
```

The file 'newthisweek.Z' was specifically included because we'll be using it later. Just for general information, it happens to be a listing of all of the

files added to UUNET's archives during the past week.

The directory shown is on a machine running the Unix operating system—the `dir` command will produce different results on other operating systems (e.g. TOPS, VMS, et al.). Learning to recognize different formats will take some time. After a few weeks of traversing the Internet, it proves easier to see, for example, how large a file is on an operating system you're otherwise not acquainted with.

With many FTP implementations, it's also possible to take the output of `dir` and put it into a file on the local system with

```
ftp> dir n* outfilename
```

the contents of which can then be read outside of the live FTP connection; this is particularly useful for systems with very long directories (like `ftp.uu.net`). The above example would put the names of every file that begins with an 'n' into the local file `outfilename`.

3.2.3 cd

At the beginning of an FTP session, the user is in a "top-level" directory. Most things are in directories below it (e.g. `'/pub'`). To change the current directory, one uses the `cd` command. To change to the directory `'pub'`, for example, one would type

```
ftp> cd pub
```

which would elicit the response

```
250 CWD command successful.
```

Meaning the "Change Working Directory" command (`'cd'`) worked properly. Moving "up" a directory is more system-specific—in Unix use the command `'cd ..'`, and in VMS, `'cd [-]'`.

3.2.4 get and put

The actual transfer is performed with the `get` and `put` commands. To get a file from the remote computer to the local system, the command takes the form:

```
ftp> get filename
```

where `filename` is the file on the remote system. Again using `ftp.uu.net` as an example, the file `'newthisweek.Z'` can be retrieved with

```
ftp> get newthisweek.Z
200 PORT command successful.
150 Opening ASCII mode data connection for newthisweek.Z (42390 bytes).
226 Transfer complete.
local: newthisweek.Z remote: newthisweek.Z
42553 bytes received in 6.9 seconds (6 Kbytes/s)
ftp>
```

The section below on using binary mode instead of ASCII will describe why this particular choice will result in a corrupt and subsequently unusable file.

If, for some reason, you want to save a file under a different name (e.g. your system can only have 14-character filenames, or can only have one dot in the name), you can specify what the local filename should be by providing `get` with an additional argument

```
ftp> get newthisweek.Z uunet-new
```

which will place the contents of the file 'newthisweek.Z' in 'uunet-new' on the local system.

The transfer works the other way, too. The `put` command will transfer a file from the local system to the remote system. If the permissions are set up for an FTP session to write to a remote directory, a file can be sent with

```
ftp> put filename
```

As with `get`, `put` will take a third argument, letting you specify a different name for the file on the remote system.

3.2.4.1 ASCII vs Binary

In the example above, the file 'newthisweek.Z' was transferred, but supposedly not correctly. The reason is this: in a normal ASCII transfer (the default), certain characters are translated between systems, to help make text files more readable. However, when *binary* files (those containing non-ASCII characters) are transferred, this translation should not take place. One example is a binary program—a few changed characters can render it completely useless.

To avoid this problem, it's possible to be in one of two modes—*ASCII* or *binary*. In *binary* mode, the file isn't translated in any way. What's on the remote system is precisely what's received. The commands to go between the two modes are:

```
ftp> ascii
200 Type set to A. (Note the A, which signifies ASCII mode.)
```

```
ftp> binary
200 Type set to I. (Set to Image format, for pure binary transfers.)
```

Note that each command need only be done once to take effect; if the user types **binary**, all transfers in that session are done in binary mode (that is, unless **ascii** is typed later).

The transfer of '**newthisweek.Z**' will work if done as:

```
ftp> binary
200 Type set to I.
ftp> get newthisweek.Z
200 PORT command successful.
150 Opening BINARY mode data connection for newthisweek.Z (42390 bytes).
226 Transfer complete.
local: newthisweek.Z remote: newthisweek.Z
42390 bytes received in 7.2 seconds (5.8 Kbytes/s)
```

Note: The file size (42390) is different from that done in ASCII mode (42553) bytes; and the number 42390 matches the one in the listing of UUNET's top directory. We can be relatively sure that we've received the file without any problems.

3.2.4.2 **mget** and **mput**

The commands **mget** and **mput** allow for multiple file transfers using wildcards to get several files, or a whole set of files at once; rather than having to do it manually one by one. For example, to get all files that begin with the letter 'f', one would type

```
ftp> mget f*
```

Similarly, to put all of the local files that end with **.c**:

```
ftp> mput *.c
```

Rather than reiterate what's been written a hundred times before, consult a local manual for more information on wildcard matching (every DOS manual, for example, has a section on it).

Normally, FTP assumes a user wants to be prompted for every file in a **mget** or **mput** operation. You'll often need to get a whole set of files and not have each of them confirmed—you know they're all right. In that case, use the **prompt** command to turn the queries off.

```
ftp> prompt
Interactive mode off.
```

Likewise, to turn it back on, the **prompt** command should simply be issued again.

3.3 The archie Server

A group of people at McGill University in Canada got together and created a query system called archie. It was originally formed to be a quick and easy way to scan the offerings of the many anonymous FTP sites that are maintained around the world. As time progressed, archie grew to include other valuable services as well.

The archie service is accessible through an interactive telnet session, email queries, and command-line and X-window clients. The email responses can be used along with FTPmail servers for those not on the Internet. (See [FTP-by-Mail Servers], page 77, for info on using FTPmail servers.)

3.3.1 Using archie Today

Currently, archie tracks the contents of over 800 anonymous FTP archive sites containing over a million files stored across the Internet. Collectively, these files represent well over 50 gigabytes of information, with new entries being added daily.

The archie server automatically updates the listing information from each site about once a month. This avoids constantly updating the databases, which could waste network resources, yet ensures that the information on each site's holdings is reasonably up to date.

To access archie interactively, **telnet** to one of the existing servers.¹ They include

- archie.ans.net (New York, USA)
- archie.rutgers.edu (New Jersey, USA)
- archie.sura.net (Maryland, USA)
- archie.unl.edu (Nebraska, USA)
- archie.mcgill.ca (the first Archie server, in Canada)
- archie.funet.fi (Finland)
- archie.au (Australia)
- archie.doc.ic.ac.uk (Great Britain)

At the **login:** prompt of one of the servers, enter 'archie' to log in. A greeting will be displayed, detailing information about ongoing work in the archie project; the user will be left at a 'archie>' prompt, at which he may enter commands. Using 'help' will yield instructions on using the 'prog' command to make queries, 'set' to control various aspects of the server's

¹ See Chapter 5 [Telnet], page 45, for notes on using the **telnet** program.

operation, et al. Type 'quit' at the prompt to leave archie. Typing the query 'prog vine.tar.Z' will yield a list of the systems that offer the source to the X-windows program vine; a piece of the information returned looks like:

```
Host ftp.uu.net (137.39.1.9)
Last updated 10:30 7 Jan 1992

Location: /packages/X/contrib
FILE      rw-r--r--  15548 Oct  8 20:29  vine.tar.Z

Host nic.funet.fi (128.214.6.100)
Last updated 05:07 4 Jan 1992

Location: /pub/X11/contrib
FILE      rw-rw-r--  15548 Nov  8 03:25  vine.tar.Z
```

3.3.2 archie Clients

There are two main-stream archie clients, one called (naturally enough) 'archie', the other 'xarchie' (for X-Windows). They query the archie databases and yield a list of systems that have the requested file(s) available for anonymous FTP, without requiring an interactive session to the server. For example, to find the same information you tried with the server command 'prog', you could type

```
% archie vine.tar.Z
Host athene.uni-paderborn.de
Location: /local/X11/more_contrib
FILE -rw-r--r--  18854 Nov 15 1990  vine.tar.Z

Host emx.utexas.edu
Location: /pub/mnt/source/games
FILE -rw-r--r--  12019 May  7 1988  vine.tar.Z

Host export.lcs.mit.edu
Location: /contrib
FILE -rw-r--r--  15548 Oct  9 00:29  vine.tar.Z
```

Note that your system administrator may not have installed the archie clients yet; the source is available on each of the archie servers, in the directory 'archie/clients'.

Using the X-windows client is much more intuitive—if it's installed, just read its man page and give it a whirl. It's essential for the networked desktop.

3.3.3 Mailing archie

Users limited to email connectivity to the Internet should send a message to the address 'archie@archie.mcgill.ca' with the single word **help** in the body of the message. An email message will be returned explaining how to use the email archie server, along with the details of using FTPmail. Most of the commands offered by the telnet interface can be used with the mail server.

3.3.4 The whatis database

In addition to offering access to anonymous FTP listings, archie also permits access to the *whatis* description database. It includes the names and brief synopses for over 3,500 public domain software packages, datasets and informational documents located on the Internet.

Additional *whatis* databases are scheduled to be added in the future. Planned offerings include listings for the names and locations of online library catalog programs, the names of publicly accessible electronic mailing lists, compilations of Frequently Asked Questions lists, and archive sites for the most popular Usenet newsgroups. Suggestions for additional descriptions or locations databases are welcomed and should be sent to the archie developers at 'archie-l@cs.mcgill.ca'.

“Was für plündern!”
("What a place to plunder!")
Gebhard Leberecht Blücher

4 Usenet News

The first thing to understand about Usenet is that it is widely misunderstood. Every day on Usenet the “blind men and the elephant” phenomenon appears, in spades. In the opinion of the author, more *flame wars* (rabid arguments) arise because of a lack of understanding of the nature of Usenet than from any other source. And consider that such flame wars arise, of necessity, among people who are on Usenet. Imagine, then, how poorly understood Usenet must be by those outside!

No essay on the nature of Usenet can ignore the erroneous impressions held by many Usenet users. Therefore, this section will treat falsehoods first. Keep reading for truth. (Beauty, alas, is not relevant to Usenet.)

4.1 What Usenet Is

Usenet is the set of machines that exchange articles tagged with one or more universally-recognized labels, called *newsgroups* (or “groups” for short). (Note that the term ‘*newsgroup*’ is correct, while ‘area’, ‘base’, ‘board’, ‘bboard’, ‘conference’, ‘round table’, ‘SIG’, etc. are incorrect. If you want to be understood, be accurate.)

4.2 The Diversity of Usenet

If the above definition of Usenet sounds vague, that’s because it is. It is almost impossible to generalize over all Usenet sites in any non-trivial way. Usenet encompasses government agencies, large universities, high schools, businesses of all sizes, home computers of all descriptions, etc.

Every administrator controls his own site. No one has any real control over any site but his own. The administrator gets his power from the owner of the system he administers. As long as the owner is happy with the job the administrator is doing, he can do whatever he pleases, up to and including cutting off Usenet entirely. *C'est la vie.*

4.3 What Usenet Is Not

Usenet is not an organization.

Usenet has no central authority. In fact, it has no central anything. There is a vague notion of “upstream” and “downstream”

related to the direction of high-volume news flow. It follows that, to the extent that "upstream" sites decide what traffic they will carry for their "downstream" neighbors, that "upstream" sites have some influence on their neighbors. But such influence is usually easy to circumvent, and heavy-handed manipulation typically results in a backlash of resentment.

Usenet is not a democracy.

A democracy can be loosely defined as "government of the people, by the people, for the people." However, as explained above, Usenet is not an organization, and only an organization can be run as a democracy. Even a democracy must be organized, for if it lacks a means of enforcing the peoples' wishes, then it may as well not exist.

Some people wish that Usenet were a democracy. Many people pretend that it is. Both groups are sadly deluded.

Usenet is not fair.

After all, who shall decide what's fair? For that matter, if someone is behaving unfairly, who's going to stop him? Neither you nor I, that's certain.

Usenet is not a right.

Some people misunderstand their local right of "freedom of speech" to mean that they have a legal right to use others' computers to say what they wish in whatever way they wish, and the owners of said computers have no right to stop them.

Those people are wrong. Freedom of speech also means freedom not to speak; if I choose not to use my computer to aid your speech, that is my right. Freedom of the press belongs to those who own one.

Usenet is not a public utility.

Some Usenet sites are publicly funded or subsidized. Most of them, by plain count, are not. There is no government monopoly on Usenet, and little or no control.

Usenet is not a commercial network.

Many Usenet sites are academic or government organizations; in fact, Usenet originated in academia. Therefore, there is a Usenet custom of keeping commercial traffic to a minimum. If such commercial traffic is generally considered worth carrying, then it may be grudgingly tolerated. Even so, it is usually separated somehow from non-commercial traffic; see `comp.newprod`.

Usenet is not the Internet.

The Internet is a wide-ranging network, parts of which are subsidized by various governments. The Internet carries many kinds of traffic; Usenet is only one of them. And the Internet is only one of the various networks carrying Usenet traffic.

Usenet is not a Unix network, nor even an ASCII network.

Don't assume that everyone is using "rn" on a Unix machine. There are Vaxes running VMS, IBM mainframes, Amigas, and MS-DOS PCs reading and posting to Usenet. And, yes, some of them use (shudder) EBCDIC. Ignore them if you like, but they're out there.

Usenet is not software.

There are dozens of software packages used at various sites to transport and read Usenet articles. So no one program or package can be called "the Usenet software."

Software designed to support Usenet traffic can be (and is) used for other kinds of communication, usually without risk of mixing the two. Such private communication networks are typically kept distinct from Usenet by the invention of newsgroup names different from the universally-recognized ones.

Usenet is not a UUCP network.

UUCP is a protocol (some might say *protocol suite*, but that's a technical point) for sending data over point-to-point connections, typically using dialup modems. Usenet is only one of the various kinds of traffic carried via UUCP, and UUCP is only one of the various transports carrying Usenet traffic.

Well, enough negativity.

4.4 Propagation of News

In the old days, when UUCP over long-distance dialup lines was the dominant means of article transmission, a few well-connected sites had real influence in determining which newsgroups would be carried where. Those sites called themselves "the backbone."

But things have changed. Nowadays, even the smallest Internet site has connectivity the likes of which the backbone admin of yesteryear could only dream. In addition, in the U.S., the advent of cheaper long-distance calls and high-speed modems has made long-distance Usenet feeds thinkable for smaller companies. There is only one pre-eminent UUCP transport site

today in the U.S., namely UUNET. But UUNET isn't a player in the propagation wars, because it never refuses any traffic—it gets paid by the minute, after all; to refuse based on content would jeopardize its legal status as an enhanced service provider.

All of the above applies to the U.S. In Europe, different cost structures favored the creation of strictly controlled hierarchical organizations with central registries. This is all very unlike the traditional mode of U.S. sites (pick a name, get the software, get a feed, you're on). Europe's "benign monopolies", long uncontested, now face competition from looser organizations patterned after the U.S. model.

4.5 Group Creation

As discussed above, Usenet is not a democracy. Nevertheless, currently the most popular way to create a new newsgroup involves a "vote" to determine popular support for (and opposition to) a proposed newsgroup. See Appendix C [Newsgroup Creation], page 79, for detailed instructions and guidelines on the process involved in making a newsgroup.

If you follow the guidelines, it is probable that your group will be created and will be widely propagated. However, due to the nature of Usenet, there is no way for any user to enforce the results of a newsgroup vote (or any other decision, for that matter). Therefore, for your new newsgroup to be propagated widely, you must not only follow the letter of the guidelines; you must also follow its spirit. And you must not allow even a whiff of shady dealings or dirty tricks to mar the vote.

So, you may ask: How is a new user supposed to know anything about the "spirit" of the guidelines? Obviously, she can't. This fact leads inexorably to the following recommendation:

If you're a new user, don't try to create a new newsgroup alone.

If you have a good newsgroup idea, then read the `news.groups` newsgroup for a while (six months, at least) to find out how things work. If you're too impatient to wait six months, then you really need to learn; read `news.groups` for a year instead. If you just can't wait, find a Usenet old hand to run the vote for you.

Readers may think this advice unnecessarily strict. Ignore it at your peril. It is embarrassing to speak before learning. It is foolish to jump into a society you don't understand with your mouth open. And it is futile to try to force your will on people who can tune you out with the press of a key.

4.6 If You're Unhappy...

Property rights being what they are, there is no higher authority on Usenet than the people who own the machines on which Usenet traffic is carried. If the owner of the machine you use says, "We will not carry `alt.sex` on this machine," and you are not happy with that order, you have no Usenet recourse. What can we outsiders do, after all?

That doesn't mean you are without options. Depending on the nature of your site, you may have some internal political recourse. Or you might find external pressure helpful. Or, with a minimal investment, you can get a feed of your own from somewhere else. Computers capable of taking Usenet feeds are down in the \$500 range now, Unix-capable boxes are going for under \$2000, and there are at least two Unix lookalikes in the \$100 price range.

No matter what, appealing to "Usenet" won't help. Even if those who read such an appeal regarding system administration are sympathetic to your cause, they will almost certainly have even less influence at your site than you do.

By the same token, if you don't like what some user at another site is doing, only the administrator and/or owner of that site have any authority to do anything about it. Persuade them that the user in question is a problem for them, and they might do something (if they feel like it). If the user in question is the administrator or owner of the site from which he or she posts, forget it; you can't win. Arrange for your newsreading software to ignore articles from him or her if you can, and chalk one up to experience.

4.7 The History of Usenet (The ABCs)

In the beginning, there were conversations, and they were good. Then came Usenet in 1979, shortly after the release of V7 Unix with UUCP; and it was better. Two Duke University grad students in North Carolina, Tom Truscott and Jim Ellis, thought of hooking computers together to exchange information with the Unix community. Steve Bellovin, a grad student at the University of North Carolina, put together the first version of the news software using shell scripts and installed it on the first two sites: `unc` and `duke`. At the beginning of 1980 the network consisted of those two sites and `phs` (another machine at Duke), and was described at the January 1980 Usenix conference in Boulder, CO.¹ Steve Bellovin later rewrote the scripts into C programs, but they were never released beyond `unc` and `duke`. Shortly

thereafter, Steve Daniel did another implementation in the C programming language for public distribution. Tom Truscott made further modifications, and this became the "A" news release.

In 1981 at the University of California at Berkeley, grad student Mark Horton and high school student Matt Glickman rewrote the news software to add functionality and to cope with the ever increasing volume of news—"A" news was intended for only a few articles per group per day. This rewrite was the "B" news version. The first public release was version 2.1 in 1982; all versions before 2.1 were considered in beta test. As The Net grew, the news software was expanded and modified. The last version maintained and released primarily by Mark was 2.10.1.

Rick Adams, then at the Center for Seismic Studies, took over coordination of the maintenance and enhancement of the news software with the 2.10.2 release in 1984. By this time, the increasing volume of news was becoming a concern, and the mechanism for moderated groups was added to the software at 2.10.2. Moderated groups were inspired by ARPA mailing lists and experience with other bulletin board systems. In late 1986, version 2.11 of news was released, including a number of changes to support a new naming structure for newsgroups, enhanced batching and compression, enhanced *ihave/sendme* control messages, and other features. The current release of news is 2.11, patchlevel 19.

A new version of news, becoming known as "C" news, has been developed at the University of Toronto by Geoff Collyer and Henry Spencer. This version is a rewrite of the lowest levels of news to increase article processing speed, decrease article expiration processing and improve the reliability of the news system through better locking, etc. The package was released to The Net in the autumn of 1987. For more information, see the paper *News Need Not Be Slow*, published in the Winter 1987 Usenix Technical Conference proceedings.

Usenet software has also been ported to a number of platforms, from the Amiga and IBM PCs all the way to minicomputers and mainframes.

4.8 Hierarchies

Newsgroups are organized according to their specific areas of concentration. Since the groups are in a tree structure, the various areas are called hierarchies. There are seven major categories:

¹ The Usenix conferences are semi-annual meetings where members of the Usenix Association, a group of Unix enthusiasts, meet and trade notes.

‘comp’	Topics of interest to both computer professionals and hobbyists, including topics in computer science, software sources, and information on hardware and software systems.
‘misc’	Group addressing themes not easily classified into any of the other headings or which incorporate themes from multiple categories. Subjects include fitness, job-hunting, law, and investments.
‘sci’	Discussions marked by special knowledge relating to research in or application of the established sciences.
‘soc’	Groups primarily addressing social issues and socializing. Included are discussions related to many different world cultures.
‘talk’	Groups largely debate-oriented and tending to feature long discussions without resolution and without appreciable amounts of generally useful information.
‘news’	Groups concerned with the news network, group maintenance, and software.
‘rec’	Groups oriented towards hobbies and recreational activities

These “world” newsgroups are (usually) circulated around the entire Usenet—this implies world-wide distribution. Not all groups actually enjoy such wide distribution, however. The European Usenet and Eunet sites take only a selected subset of the more “technical” groups, and controversial “noise” groups are often not carried by many sites in the U.S. and Canada (these groups are primarily under the ‘talk’ and ‘soc’ classifications). Many sites do not carry some or all of the `comp.binaries` groups because of the typically large size of the posts in them (being actual executable programs).

Also available are a number of “alternative” hierarchies:

‘alt’	True anarchy; anything and everything can and does appear; subjects include sex, the Simpsons, and privacy.
‘gnu’	Groups concentrating on interests and software with the GNU Project of the Free Software Foundation. For further info on what the FSF is, see Section 8.3.4 [FSF], page 68.
‘biz’	Business-related groups.

4.9 Moderated vs Unmoderated

Some newsgroups insist that the discussion remain focused and on-target; to serve this need, moderated groups came to be. All articles posted to a moderated group get mailed to the group’s *moderator*. He or she periodically (hopefully sooner than later) reviews the posts, and then either posts them

individually to Usenet, or posts a composite digest of the articles for the past day or two. This is how many mailing list gateways work (for example, the *Risks Digest*).

4.10 news.groups & news.announce.newgroups

Being a good *net.citizen* includes being involved in the continuing growth and evolution of the Usenet system. One part of this involvement includes following the discussion in the groups **news.groups** and the notes in **news.announce.newgroups**. It is there that discussion goes on about the creation of new groups and destruction of inactive ones. Every person on Usenet is allowed and encouraged to vote on the creation of a newsgroup.

4.11 How Usenet Works

The transmission of Usenet news is entirely cooperative. Feeds are generally provided out of good will and the desire to distribute news everywhere. There are places which provide feeds for a fee (e.g. UUNET), but for the large part no exchange of money is involved.

There are two major transport methods, UUCP and NNTP. The first is mainly modem-based and involves the normal charges for telephone calls. The second, NNTP, is the primary method for distributing news over the Internet.

With UUCP, news is stored in *batches* on a site until the neighbor calls to receive the articles, or the feed site happens to call. A list of groups which the neighbor wishes to receive is maintained on the feed site. The Cnews system compresses its batches, which can dramatically reduce the transmission time necessary for a relatively heavy newsfeed.

NNTP, on the other hand, offers a little more latitude with how news is sent. The traditional store-and-forward method is, of course, available. Given the "real-time" nature of the Internet, though, other methods have been devised. Programs now keep constant connections with their news neighbors, sending news nearly instantaneously, and can handle dozens of simultaneous feeds, both incoming and outgoing.

The transmission of a Usenet article is centered around the unique 'Message-ID:' header. When an NNTP site offers an article to a neighbor, it says it has that specific Message ID. If the neighbor finds it hasn't received the article yet, it tells the feed to send it through; this is repeated for each and every article that's waiting for the neighbor. Using unique IDs

helps prevent a system from receiving five copies of an article from each of its five news neighbors, for example.

Further information on how Usenet works with relation to the various transports is available in the documentation for the Cnews and NNTP packages, as well as in RFC-1036, the *Standard for Interchange of USENET Messages* and RFC-977, *Network News Transfer Protocol: A Proposed Standard for the Stream-Based Transmission of News*. The RFCs do tend to be rather dry reading, particularly to the new user. See [RFCs], page 73 for information on retrieving RFCs.

4.12 Mail Gateways

A natural progression is for Usenet news and electronic mailing lists to somehow become merged—which they have, in the form of *news gateways*. Many mailing lists are set up to “reflect” messages not only to the readership of the list, but also into a newsgroup. Likewise, posts to a newsgroup can be sent to the moderator of the mailing list, or to the entire mailing list. Some examples of this in action are `comp.risks` (the *Risks Digest*) and `comp.dcom.telecom` (the *Telecom Digest*).

This method of propagating mailing list traffic has helped solve the problem of a single message being delivered to a number of people at the same site—instead, anyone can just subscribe to the group. Also, mailing list maintenance is lowered substantially, since the moderators don’t have to be constantly removing and adding users to and from the list. Instead, the people can read and not read the newsgroup at their leisure.

4.13 Usenet “Netiquette”

There are many traditions with Usenet, not the least of which is dubbed *netiquette*—being polite and considerate of others. If you follow a few basic guidelines, you, and everyone that reads your posts, will be much happier in the long run.

4.13.1 Signatures

At the end of most articles is a small blurb called a person’s *signature*. In Unix this file is named ‘`.signature`’ in the person’s login directory—it will vary for other operating systems. It exists to provide information about how to get in touch with the person posting the article, including

their email address, phone number, address, or where they're located. Even so, signatures have become the graffiti of computers. People put song lyrics, pictures, philosophical quotes, even advertisements in their ".sigs". (Note, however, that advertising in your signature will more often than not get you flamed until you take it out.)

Four lines will suffice—more is just extra garbage for Usenet sites to carry along with your article, which is supposed to be the intended focus of the reader. Netiquette dictates limiting oneself to this "quota" of four—some people make signatures that are ten lines or even more, including elaborate ASCII drawings of their hand-written signature or faces or even the space shuttle. This is *not* cute, and will bother people to no end.

Similarly, it's not necessary to include your signature—if you forget to append it to an article, *don't worry* about it. The article's just as good as it ever would be, and contains everything you should want to say. Don't re-post the article just to include the signature.

4.13.2 Posting Personal Messages

If mail to a person doesn't make it through, avoid posting the message to a newsgroup. Even if the likelihood of that person reading the group is very high, all of the other people reading the articles don't give a whit what you have to say to Jim Morrison. Simply wait for the person to post again and double-check the address, or get in touch with your system administrator and see if it's a problem with local email delivery. It may also turn out that their site is down or is having problems, in which case it's just necessary to wait until things return to normal before contacting Jim.

4.13.3 Posting Mail

In the interests of privacy, it's considered extremely bad taste to post any email that someone may have sent, unless they explicitly give you permission to redistribute it. While the legal issues can be heavily debated, most everyone agrees that email should be treated as anything one would receive via normal snailmail,², with all of the assumed rights that are carried with it.

² The slang for the normal land and air postal service.

4.13.4 Test Messages

Many people, particularly new users, want to try out posting before actually taking part in discussions. Often the mechanics of getting messages out is the most difficult part of Usenet. To this end, many, many users find it necessary to post their tests to "normal" groups (for example, `news.admin` or `comp.mail.misc`). This is considered a major netiquette *faux pas* in the Usenet world. There are a number of groups available, called *test groups*, that exist solely for the purpose of trying out a news system, reader, or even new signature. They include

```
alt.test
gnu.gnusenet.test
misc.test
```

some of which will generate auto-magic replies to your posts to let you know they made it through. There are certain denizens of Usenet that frequent the test groups to help new users out. They respond to the posts, often including the article so the poster can see how it got to the person's site. Also, many regional hierarchies have test groups, like `phl.test` in Philadelphia.

By all means, experiment and test—just do it in its proper place.

4.13.5 Famous People Appearing

Every once in a while, someone says that a celebrity is accessible through "The Net"; or, even more entertaining, an article is forged to appear to be coming from that celebrity. One example is Stephen Spielberg—the `rec.arts.movies` readership was in an uproar for two weeks following a couple of posts supposedly made by Mr. Spielberg. (Some detective work revealed it to be a hoax.)

There are a few well-known people that are acquainted with Usenet and computers in general—but the overwhelming majority are just normal people. One should act with skepticism whenever a notable personality is "seen" in a newsgroup.

4.13.6 Summaries

Authors of articles occasionally say that readers should reply by mail and they'll summarize. Accordingly, readers should do just that—reply via mail. Responding with a followup article to such an article defeats the intention of the author. She, in a few days, will post one article containing the highlights of the responses she received. By following up to the whole group, the author may not read what you have to say.

When creating a summary of the replies to a post, try to make it as *reader-friendly* as possible. Avoid just putting all of the messages received into one big file. Rather, take some time and edit the messages into a form that contains the essential information that other readers would be interested in.

Also, sometimes people will respond but request to remain anonymous (one example is the employees of a corporation that feel the information's not proprietary, but at the same time want to protect themselves from political backlash). Summaries should honor this request accordingly by listing the 'From:' address as 'anonymous' or '(Address withheld by request)'.

4.13.7 Quoting

When following up to an article, many newsreaders provide the facility to quote the original article with each line prefixed by '>', as in

```
In article <1232@foo.bar.com>, sharon@foo.bar.com wrote:  
> I agree, I think that basketweaving's really catching on,  
> particularly in Pennsylvania. Here's a list of every person  
> in PA that currently engages in it publicly:  
  
... etc ...
```

This is a severe example (potentially a horribly long article), but proves a point. When you quote another person, *edit out* whatever isn't directly applicable to your reply.³ This gives the reader of the new article a better idea of what points you were addressing. By including the entire article, you'll only annoy those reading it. Also, signatures in the original aren't necessary; the readers already know who wrote it (by the attribution).

Avoid being tedious with responses—rather than pick apart an article, address it in parts or as a whole. Addressing practically each and every word in an article only proves that the person responding has absolutely nothing better to do with his time.

If a "war" starts (insults and personal comments get thrown back and forth), *take it into email*—exchange email with the person you're arguing with. No one enjoys watching people bicker incessantly.

³ But not changing their words, of course.

4.13.8 Crossposting

The 'Newsgroups:' line isn't limited to just one group—an article can be posted in a list of groups. For instance, the line

Newsgroups: sci.space,comp.simulation

posts the article to both the groups `sci.space` and `comp.simulation`. It's usually safe to crosspost to up to three or four groups. To list more than that is considered "excessive noise."

It's also suggested that if an article is crossposted a 'Followup-To:' header be included. It should name the group to which all additional discussion should be directed to. For the above example a possible 'Followup-To:' would be

Followup-To: sci.space

which would make all followups automatically be posted to just `sci.space`, rather than both `sci.space` and `comp.simulation`. If every response made with a newsreader's "followup" command should go to the person posting the article no matter what, there's also a mechanism worked in to accommodate. The `Followup-To:` header should contain the single word 'poster':

Followup-To: poster

Certain newsreaders will use this to sense that a reply should never be posted back onto The Net. This is often used with questions that will yield a summary of information later, a vote, or an advertisement.

4.13.9 Recent News

One should avoid posting "recent" events—sports scores, a plane crash, or whatever people will see on the evening news or read in the morning paper. By the time the article has propagated across all of Usenet, the "news" value of the article will have become stale. (This is one case for the argument that 'Usenet news' is a misnomer.⁴)

4.13.10 Quality of Postings

How you write and present yourself in your articles is important. If you have terrible spelling, keep a dictionary near by. If you have trouble

⁴ Note that the Clarinet News service (see Section 7.3 [Clarinet], page 60) offers news items in a Usenet format as a precise alternative to the morning paper, et. al.

with grammar and punctuation, try to get a book on English grammar and composition (found in many bookstores and at garage sales). By all means pay attention to what you say—it makes you who you are on The Net.

Likewise, try to be clear in what you ask. Ambiguous or vague questions often lead to no response at all, leaving the poster discouraged. Give as much essential information as you feel is necessary to let people help you, but keep it within limits. For instance, you should probably include the operating system of your computer in the post if it's needed, but don't tell everybody what peripherals you have hanging off of it.

4.13.11 Useful Subjects

The 'Subject:' line of an article is what will first attract people to read it—if it's vague or doesn't describe what's contained within, no one will read the article. At the same time, 'Subject:' lines that're too wordy tend to be irritating. For example:

Good	Subject: Building Emacs on a Sun Sparc under 4.1
Good	Subject: Tryin' to find Waldo in NJ.
Bad	Subject: I can't get emacs to work !!!
Bad	Subject: I'm desperately in search of the honorable Mr. Waldo in the state of...

Simply put, try to think of what will best help the reader when he or she encounters your article in a newsreading session.

4.13.12 Tone of Voice

Since common computers can't portray the inflection or tone in a person's voice, how articles are worded can directly affect the response to them. If you say

Anybody using a Vic-20 should go buy themselves a life.
you'll definitely get some responses—telling you to take a leap. Rather than be inflammatory, phrase your articles in a way that rationally expresses your opinion, like

What're the practical uses of a Vic-20 these days?
which presents yourself as a much more level-headed individual.

Also, what case (upper or lower) you use can indicate how you're trying to speak—netiquette dictates that if you USE ALL CAPITAL LETTERS, people will think you're "shouting." Write as you would in a normal letter

to a friend, following traditional rules of English (or whatever language you happen to speak).

4.13.13 Computer Religion

No matter what kind of computer a person is using, theirs is always the best and most efficient of them all. Posting articles asking questions like 'What computer should I buy? An Atari ST or an Amiga?' will lead only to fervent arguments over the merits and drawbacks of each brand. Don't even ask The Net—go to a local user group, or do some research of your own like reading some magazine reviews. Trying to say one computer is somehow better than another is a moot point.

4.14 Frequently Asked Questions

A number of groups include *Frequently Asked Question* (FAQ) lists, which give the answers to questions or points that have been raised time and time again in a newsgroup. They're intended to help cut down on the redundant traffic in a group. For example, in the newsgroup `alt.tv.simpsons`, one recurring question is 'Did you notice that there's a different blackboard opening at the beginning of every *Simpsons* episode?' As a result, it's part of the FAQ for that group.

Usually, FAQ lists are posted at the beginning of each month, and are set to expire one month later (when, supposedly, the next FAQ will be published). Nearly every FAQ is also crossposted to `news.answers`, which is used as a Usenet repository for them.

4.14.1 The Pit-Manager Archive

MIT, with Jonathan Kamens, has graciously dedicated a machine to the archiving and storage of the various periodic postings that are peppered throughout the various Usenet groups. To access them, FTP to the system `pit-manager.mit.edu` and look in the directory '`/pub/usenet`'.

"Be it true or false, so it be news."
Ben Jonson, *News from the New World*

5 Telnet

Telnet is the main Internet protocol for creating a connection with a remote machine. It gives the user the opportunity to be on one computer system and do work on another, which may be across the street or thousands of miles away. Where modems are limited, in the majority, by the quality of telephone lines and a single connection, telnet provides a connection that's error-free and nearly always faster than the latest conventional modems.

5.1 Using Telnet

As with FTP (see Section 3.2.2 [Anonymous FTP], page 21), the actual command for negotiating a telnet connection varies from system to system. The most common is `telnet` itself, though. It takes the form of:

```
telnet somewhere.domain
```

To be safe, we'll use your local system as a working example. By now, you hopefully know your site's domain name. If not, ask or try to figure it out. You'll not get by without it.

To open the connection, type

```
telnet your.system.name
```

If the system were `wubba.cs.widener.edu`, for example, the command would look like

```
telnet wubba.cs.widener.edu
```

The system will respond with something similar to

```
Trying 147.31.254.999...
Connected to wubba.cs.widener.edu.
Escape character is '^]'.

```

The escape character, in this example `^]` (Control-`]`), is the character that will let you go back to the local system to close the connection, suspend it, etc. To close this connection, the user would type `^]`, and respond to the `telnet>` prompt with the command `close`. Local documentation should be checked for information on specific commands, functions, and escape character that can be used.

5.1.1 Telnet Ports

Many telnet clients also include a third option, the *port* on which the connection should take place. Normally, port 23 is the default telnet port;

the user never has to think about it. But sometimes it's desirable to telnet to a different port on a system, where there may be a service available, or to aid in debugging a problem. Using

```
telnet somewhere.domain port
```

will connect the user to the given *port* on the system *somewhere.domain*. Many libraries use this port method to offer their facilities to the general Internet community; other services are also available. For instance, one would type

```
telnet martini.eecs.umich.edu 3000
```

to connect to the geographic server at the University of Michigan (see Section 5.5.6 [Geographic Server], page 50). Other such port connections follow the same usage.

5.2 Publicly Accessible Libraries

Over the last several years, most university libraries have switched from a manual (card) catalog system to computerized library catalogs. The automated systems provide users with easily accessible and up-to-date information about the books available in these libraries. This has been further improved upon with the advent of local area networks, dialup modems, and wide area networks. Now many of us can check on our local library's holdings or that of a library halfway around the world!

Many, many institutions of higher learning have made their library catalogs available for searching by anyone on the Internet. They include Boston University, the Colorado Alliance of Research Libraries (CARL), and London University King's College.

To include a listing of some of the existing sites would not only be far too long for this document, it would soon be out of date. Instead, several lists are being maintained and are available either by mail or via FTP. Also, the Internet Resource Guide (IRG) also describes a few libraries that are accessible—see Section 9.1 [IRG], page 71 for further information.

Art St. George and Ron Larsen are maintaining a list of Internet-accessible libraries and databases often referred to as "the St. George directory." It began with only library catalogs but has expanded to include sections on campus-wide information systems, and even bulletin board systems that are not on the Internet. The library catalog sections are divided into those that are free, those that charge, and international (i.e. non-U.S.) catalogs; they are arranged by state, province, or country within each section. There is also a section giving dialup information for some of the library catalogs. It's available for FTP (see

Section 3.2.2 [Anonymous FTP], page 21) on `nic.cerf.net` in the directory '`cerfnet/cerfnet_info/library_catalog`'. The file '`internet-catalogs`' has a date suffix; check for the most current date. The information is updated periodically.

Billy Barron, Systems Manager at the University of North Texas, produces a directory as an aid to his user community. It complements the St. George guide by providing a standard format for all systems which lists the Internet address, login instructions, the system vendor, and logoff information. The arrangement is alphabetic by organization name. It's available for FTP on `vaxb.acs.unt.edu` in the subdirectory '`library`' as the file '`libraries.txt`'.

For announcements of new libraries being available and discussion on related topics, consult the Usenet newsgroup `comp.internet.library` (see Chapter 4 [Usenet News], page 29 to learn how to read news).

5.3 The Cleveland Freenet

Freenets are open-access, free, community computer systems. One such system is the Cleveland Freenet, sponsored by CWRU (Case Western Reserve University). Anyone and everyone is welcome to join and take part in the exciting project—that of a National Telecomputing Public Network, where everyone benefits. There's *no* charge for the registration process and *no* charge to use the system.

To register, telnet to any one of

`freenet-in-a.cwru.edu`
`freenet-in-b.cwru.edu`
`freenet-in-c.cwru.edu`

After you're connected, choose the entry on the menu that signifies you're a guest user. Another menu will follow; select '*Apply for an account*', and you'll be well on your way to being a FreeNet member.

You will need to fill out a form and send it to them through the Postal Service—your login id and password will be created in a few days. At that point you're free to use the system as you wish. They provide multi-user chat, email, Usenet news, and a variety of other things to keep you occupied for hours on end.

5.4 Directories

There are a few systems that are maintained to provide the Internet com-

munity with access to lists of information—users, organizations, etc. They range from fully dedicated computers with access to papers and research results, to a system to find out about the faculty members of a university.

5.4.1 Knowbot

Knowbot is a “master directory” that contains email address information from the NIC WHOIS database (see Section 6.4.1 [Whois], page 57), the PSI White Pages Pilot Project, the NYSERNET X.500 database and MCI Mail. Most of these services are email registries themselves, but Knowbot provides a very comfortable way to access all of them in one place. Telnet to `nri.reston.va.us` on port 185.

5.4.2 White Pages

PSI maintains a directory of information on individuals. It will list the person’s name, organization, and email address if it is given. Telnet to `wp.psi.net` and log in as ‘fred’. The White Pages Project also includes an interface to use Xwindows remotely.

5.5 Databases

For information on database services, see Section 7.2 [Commercial Databases], page 60. Not all databases on the Internet require payment for use, though. There do exist some, largely research-driven databases, which are publicly accessible. New ones spring up regularly.

To find out more about the databases in this section, contact the people directly responsible for them. Their areas of concentration and the software used to implement them are widely disparate, and are probably beyond the author’s expertise. Also, don’t forget to check with your local library—the reference librarian there can provide information on conventional resources, and possibly even those available over the Internet (they are becoming more common).

5.5.1 Colorado Alliance of Research Libraries (CARL)

The Colorado Alliance of Research Libraries (CARL), in association with CARL Systems Inc., operates a public access catalog of services. Offered are a number of library databases, including searches for government periodicals, book reviews, indices for current articles, and access to to other library

databases around the country. Other services are available to CARL members including an online encyclopedia. Telnet to `pac.carl.org`, or write to '`help@carl.org`' for more details.

5.5.2 PENpages

PENpages is an agriculturally-oriented database administered by Pennsylvania State University. Information entered into PENpages is provided by numerous sources including the Pennsylvania Dept. of Agriculture, Rutgers University, and Penn State. Easy-to-use menus guide users to information ranging from cattle and agricultural prices to current weather information, from health information to agricultural news from around the nation. A keyword search option also allows users to search the database for related information and articles. The database is updated daily, and a listing of most recent additions is displayed after login. Telnet to `psupen.psu.edu` and log in as the user 'PNOTPA'.

5.5.3 Clemson Univ. Forestry & Agricultural Network

Clemson maintains a database similar to PENpages in content, but the information provided tends to be localized to the Southeastern United States. A menu-driven database offers queries involving the weather, food, family, and human resources. Telnet to `eureka.clemson.edu` and log in as 'PUBLIC'. You need to be on a good VT100 emulator (or a real VT terminal).

5.5.4 University of Maryland Info Database

The Computer Science department of the University of Maryland maintains a repository of information on a wide variety of topics. They wish to give a working example of how network technology can (and should) provide as much information as possible to those who use it. Telnet to `info.umd.edu` and log in as 'info'. The information contained in the database is accessible through a screen-oriented interface, and everything therein is available via anonymous FTP.

There is a mailing list used to discuss the UMD Info Database, welcoming suggestions for new information, comments on the interface the system provides, and other related topics. Send mail to `listserv@umdd.umd.edu` with a body of

`subscribe INFO-L Your Full Name`

See Section 2.2.1 [Listservs], page 16 for more information on using the Listserv system.

5.5.5 University of Michigan Weather Underground

The University of Michigan's Department of Atmospheric, Oceanic, & Space Sciences maintains a database of weather and related information for the United States and Canada. Available are current weather conditions and forecasts for cities in the U.S., a national weather summary, ski conditions, earthquake and hurricane updates, and a listing of severe weather conditions. Telnet to `madlab.spqr.umich.edu` on port 3000 to use the system.

5.5.6 Geographic Name Server

A geographic database listing information for cities in the United States and some international locations is maintained by Merit, Inc. The database is searchable by city name, zip code, etc. It will respond with a lot of information: the area code, elevation, time zone, and longitude and latitude are included. For example, a query of '19013' yields

```
0 Chester
1 42045 Delaware
2 PA Pennsylvania
3 US United States
F 45 Populated place
L 39 50 58 N 75 21 22 W
P 45794
E 22
Z 19013
Z 19014
Z 19015
Z 19016
```

To use the server, telnet to `martini.eecs.umich.edu` on port 3000. The command 'help' will yield further instructions, along with an explanation for each of the fields in a response.

5.5.7 FEDIX—Minority Scholarship Information

FEDIX is an on-line information service that links the higher education community and the federal government to facilitate research, education, and services. The system provides accurate and timely federal agency information to colleges, universities, and other research organizations. There are no registration fees and no access charges for FEDIX whatsoever.

FEDIX offers the Minority On-Line Information Service (MOLIS), a database listing current information about Black and Hispanic colleges and universities.

Daily information updates are made on federal education and research programs, scholarships, fellowships, and grants, available used research equipment, and general information about FEDIX itself. To access the database, telnet to fedix.fie.com and log in as 'fedix'.

5.5.8 Science & Technology Information System

The STIS is maintained by the National Science Foundation (NSF), and provides access to many NSF publications. The full text of publications can be searched online and copied from the system, which can accommodate up to ten users at one time. Telnet to stis.nsf.gov and log in as 'public'. Everything on the system is also available via anonymous FTP. For further information, contact:

STIS, Office of Information Systems, Room 401
National Science Foundation
1800 G. Street, N.W.
Washington, D.C. 20550
stis-request@nsf.gov
(202) 357-7492
(202) 357-7663 (Fax)

5.5.9 Ocean Network Information Center

The University of Delaware College of Marine Studies offers access to an interactive database of research information covering all aspects of marine studies, nicknamed OCEANIC. This includes the World Oceanic Circulation Experiment (WOCE) information and program information, research ship schedules and information, and a Who's Who of email and mailing addresses for oceanic studies. Data from a variety of academic institutions based on research studies is also available. Telnet to delocn.udel.edu and log in as 'INFO'.

5.5.10 NASA/IPAC Extragalactic Database (NED)

The NASA/IPAC Extragalactic Database (NED) is an ongoing project, funded by NASA, to make data and literature on extragalactic objects available over computer networks. NED is an object-oriented database which

contains extensive information for nearly 132,000 extragalactic objects taken from about major catalogs of galaxies, quasars, infrared and radio sources. NED provides positions, names, and other basic data (e.g. magnitude types, sizes and redshifts as well as bibliographic references and abstracts). Searches can be done by name, around a name, and on an astronomical position. NED contains a tutorial which guides the user through the retrieval process. Telnet to `ipac.caltech.edu` and log in as 'ned'.

5.5.11 U.S. Naval Observatory Automated Data Service

Operated by the U.S. Naval Observatory in Washington, D.C., this automated data service provides database access to information ranging from current navigational satellite positioning, astronomical data, and software utilities. A wide variety of databases can be searched and instructions for file transfer are given. Telnet to `tycho.usno.navy.mil` and log in as 'ads'.

“My consciousness suddenly switched locations, for the first time in my life, from the vicinity of my head and body to a point about twenty feet away from where I normally see the world.”

Howard Rheingold, *Virtual Reality*

6 Various Tools

New and interesting ways to use the Internet are being dreamed up every day. As they gain wide-spread use, some methods become near-standard (or actual written standard) tools for Internet users to take advantage of. A few are detailed here; there are undoubtedly others, and new ideas spring up all the time. An active user of the Internet will discover most of the more common ones in time. Usually, these services are free. See Chapter 7 [Commercial Services], page 59 for applications that are commercially available over the Internet.

Usenet is often used to announce a new service or capability on the Internet. In particular, the groups `comp.archives` and `comp.protocols.tcp-ip` are good places to look. Information will drift into other areas as word spreads. See Chapter 4 [Usenet News], page 29 for information on reading news.

6.1 Finger

On many systems there exists the ‘`finger`’ command, which yield information about each user that’s currently logged in. This command also has extensions for use over the Internet, as well. Under normal circumstances, the command is simply ‘`finger`’ for a summary of who’s logged into the local system, or ‘`finger username`’ for specific information about a user. It’s also possible to go one step further and go onto the network. The general usage is

```
finger @hostname
```

To see who’s currently logged in at Widener University, for instance, use

```
% finger @cs.widener.edu
[cs.widener.edu]
Login      Name          TTY  Idle    When          Where
brendan    Brendan Kehoe    p0    Fri 02:14  tattoo.cs.widene
sven       Sven Heinicke   p1    Fri 04:16  xplex3.cs.widen
```

To find out about a certain user, they can be fingered specifically (and need not be logged in):

```
% finger bart@cs.widener.edu
[cs.widener.edu]
Login name: bart                                In real life: Bart Simpson
Directory: /home/springfield/bart                Shell: /bin/underachiever
Affiliation: Brother of Lisa                     Home System: channel29.fox.org
Last login Thu May 23 12:14 (EDT) on ttym6 from channel29.fox.org.
No unread mail
Project: To become a "fluff" cartoon character.
Plan:
Don't have a cow, man.
```

Please realize that some sites are very security conscious, and need to restrict the information about their systems and users available to the outside world. To that end, they often block finger requests from outside sites—so don't be surprised if fingering a computer or a user returns with 'Connection refused'.

6.2 Ping

The 'ping' command allows the user to check if another system is currently "up" and running. The general form of the command is 'ping system'.¹ For example,

```
ping cs.widener.edu
```

will tell you if the main machine in Widener University's Computer Science lab is currently online (we certainly hope so!).

Many implementations of 'ping' also include an option to let you see how fast a link is running (to give you some idea of the load on the network). For example:

```
% ping -s cs.swarthmore.edu
PING cs.swarthmore.edu: 56 data bytes
64 bytes from 130.58.68.1: icmp_seq=0 ttl=251 time=66 ms
64 bytes from 130.58.68.1: icmp_seq=1 ttl=251 time=45 ms
64 bytes from 130.58.68.1: icmp_seq=2 ttl=251 time=46 ms
^C
--- cs.swarthmore.edu ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 45/52/66 ms
```

This case tells us that for 'cs.swarthmore.edu' it takes about 46 milliseconds for a packet to go from Widener to Swarthmore College and back again.

¹ The usage will, again, vary.

It also gives the average and worst-case speeds, and any packet loss that may have occurred (e.g. because of network congestion).

While ‘ping’ generally doesn’t hurt network performance, you shouldn’t use it too often—usually once or twice will leave you relatively sure of the other system’s state.

6.3 Talk

Sometimes email is clumsy and difficult to manage when one really needs to have an interactive conversation. The Internet provides for that as well, in the form of *talk*. Two users can literally see each other type across thousands of miles.

To talk with Bart Simpson at Widener, one would type

```
talk bart@cs.widener.edu
```

which would cause a message similar to the following to be displayed on Bart’s terminal:

```
Message from Talk_Daemon@cs.widener.edu at 21:45 ...
talk: connection requested by joe@ee.someplace.edu
talk: respond with: talk joe@ee.someplace.edu
```

Bart would, presumably, respond by typing ‘talk joe@ee.someplace.edu’. They could then chat about whatever they wished, with instantaneous response time, rather than the write-and-wait style of email. To leave *talk*, on many systems one would type **Ctrl-C** (hold down the Control key and press ‘C’). Check local documentation to be sure.

There are two different versions of *talk* in common use today. The first, dubbed “old talk,” is supported by a set of Unix systems (most notably, those currently sold by Sun). The second, *ntalk* (aka “new talk”), is more of the standard. If, when attempting to talk with another user, it responds with an error about protocol families, odds are the incompatibilities between versions of *talk* is the culprit. It’s up to the system administrators of sites which use the old *talk* to install *ntalk* for their users.

6.4 The WHOIS Database

The main *WHOIS* database is run at the Network Information Center (NIC). The ‘*whois*’ command will let you search a database of every registered domain (e.g. ‘*mit.edu*’) and of registered users. It’s primarily used by system postmasters or listowners to find the *Points of Contact* for a site, to

let them know of a problem or contact them for one reason or another. You can also find out their postal address. For example:

```
% whois mit.edu
Massachusetts Institute of Technology (MIT) MIT.EDU          18.72.2.1
Massachusetts Institute of Technology (MIT-DOM)          MIT.EDU
```

Note that there are two entries for 'mit.edu'; we'll go for the second.

```
% whois mit-dom
Massachusetts Institute of Technology (MIT-DOM) ⇒ Mailing address
Cambridge, MA 02139

Domain Name: MIT.EDU ⇒ Domain name

Administrative Contact, Technical Contact, Zone Contact:
Schiller, Jeffrey I. (JIS) JIS@MIT.EDU
(617) 253-8400

Record last updated on 22-Jun-88. ⇒ Last change made to the record

Domain servers in listed order: ⇒ Systems that can tell you the Internet
                                addresses for a site
STRAWB/MIT.EDU          18.71.0.151
W20NS/MIT.EDU          18.70.0.160
BITSY/MIT.EDU          18.72.0.3
LITHIUM.LCS/MIT.EDU      18.26.0.121
```

To see this host record with registered users, repeat the command with a star ('*') before the name; or, use '%' to show JUST the registered users.

Much better! Now this information (sought, possibly, by a system administrator) can be used to find out how to notify MIT of a security issue or problem with connectivity.

Queries can be made for individuals as well; the following would yield an entry for the author:

```
% whois brendan
Kehoe, Brendan (BK59) brendan@cs.widener.edu
Widener University
Department of Computer Science
Kirkbride 219
P.O. Box 83 Widener University
Chester, PA 19013
(215)/499-4011
```

Record last updated on 02-May-91.

Included is the author's name, his *handle* (a unique sequence of letters and numbers), information on how to contact him, and the last time the record was modified in any way.

Anyone can register with the whois database. People who are administrative or technical contacts for domains are registered automatically when their domain applications are processed. For normal users, one must simply fill out a form from the NIC. FTP to `nic.ddn.mil` and get the file '`netinfo/user-template.txt`'. The completed form should be mailed to '`registrar@nic.ddn.mil`'.

6.4.1 Other Uses of WHOIS

Also, many educational sites run WHOIS servers of their own, to offer information about people who may be currently on the staff or attending the institution. To specify a WHOIS server, many implementations include some sort of option or qualifier—in VMS under MultiNet, it's '`/HOST`', in Unix '`-h`'. To receive information about using the Stanford server, one might use the command

```
whois -h stanford.edu help
```

A large list of systems offering WHOIS services is being maintained by Matt Power of MIT (`mhpower@stan.mit.edu`). It is available via anonymous FTP from `sipb.mit.edu`, in the directory '`pub/whois`'. The file is named '`whois-servers.list`'.

The systems available include, but are certainly not limited to, Syracuse University (`syr.edu`), New York University (`acfcluster.nyu.edu`), the University of California at San Diego (`ucsd.edu`), and Stanford University (`stanford.edu`).

"Fingers were made before forks."
Jonathan Swift, *Polite Conversation*

7 Commercial Services

Many services can be accessed through the Internet. As time progresses and more outlets for commercial activity appear, once-restricted traffic (by the NSFnet Acceptable Use Policy) may now flow freely. Now that there are other networks for that information to travel on, businesses are making their move.

7.1 Electronic Journals

The Association of Research Libraries (ARL) publishes a hard-copy directory of electronic journals, newsletters, and scholarly discussion lists. It is a compilation of entries for hundreds of scholarly lists, dozens of journals and newsletters, and a many "other" titles, including newsletter-digests, into one reference source. Each entry includes instructions on how to access the referenced publication or list.

The documents are available electronically by sending the commands

```
get ejournl1 directry  
get ejournl2 directry
```

to the server at 'LISTSERV@OTTAWA.BITNET'. See Section 2.2.1 [Listservs], page 16 for further instructions on using a listserv.

The directory, along with a compilation by Diane Kovacs called *Directories of Academic E-Mail Conferences*, is available in print and on diskette (DOS WordPerfect and MacWord) from:

Office of Scientific & Academic Publishing
Association of Research Libraries
1527 New Hampshire Avenue, NW
Washington, DC 20036
ARLHQ@UMDC.BITNET
(202) 232-2466
(202) 462-7849 (Fax)

The ARL is a not-for-profit organization representing over one hundred research libraries in the United States and Canada. The publication is available to ARL members for \$10 and to non-members for \$20 (add \$5 postage per directory for foreign addresses). Orders of six or more copies will receive a 10% discount; all orders must be prepaid and sent to the ARL.

7.2 Commercial Databases

The American Institute of Physics maintains the Physics Information Network. It contains the bibliographic SPIN and General Physics Advanced Abstracts databases. Also available is access to bulletin boards and several searchable lists (job notices, announcements, etc). Telnet to [pinet.aip.org](telnet://pinet.aip.org); new users must log in as 'NEW' and give registration information.

7.3 Clarinet News

Clarinet's an electronic publishing network service that provides professional news and information, including live UPI wireservice news, in the Usenet file format. See Chapter 4 [Usenet News], page 29 for more information about Usenet.

Clarinet lets you read an "electronic newspaper" right on the local system; you can get timely industry news, technology related wirestories, syndicated columns and features, financial information, stock quotes and more.

Clarinet's provided by using the Usenet message interchange format, and is available via UUCP and other delivery protocols, including NNTP.

The main feature is *ClariNews*, an "electronic newspaper," gathered live from the wire services of United Press International (UPI). *ClariNews* articles are distributed in 100 newsgroups based on their subject matter, and are keyworded for additional topics and the geographical location of the story. *ClariNews* includes headlines, industry news, box scores, network TV schedules, and more. The main products of *ClariNews* are:

- **ClariNews General**, the general news "paper" with news, sports, and features, averaging about 400 stories per day.
- **TechWire**, special groups for stories on science, technology, and industry stories around them.
- **ClariNews-Biz**, business and financial stories.
- **Newsbytes**, a daily computer industry newsmagazine.
- **Syndicated Columns**, including Dave Barry (humor) and Mike Royko (opinion).

Full information on ClariNet, including subscription information, is available from

Clarinet Communications Corp.
124 King St. North
Waterloo, Ontario N2J 2X8
info@clarinet.com

(800) USE-NETS

or with anonymous FTP in the directory '/Clarinet' on `ftp.uu.net` (see Section 3.2.2 [Anonymous FTP], page 21).

“Needless to say, Aristotle did not envisage modern finance.”

Frederick Copleston, S.J.
A History of Philosophy, v.1

8 Things You'll Hear About

There are certain things that you'll hear about shortly after you start actively using the Internet. Most people assume that everyone's familiar with them, and they require no additional explanation. If only that were true!

This section addresses a few topics that are commonly encountered and asked about as a new user explores Cyberspace. Some of them are directly related to how the networks are run today; other points are simply interesting to read about.

8.1 The Internet Worm

On November 2, 1988, Robert Morris, Jr., a graduate student in Computer Science at Cornell, wrote an experimental, self-replicating, self-propagating program called a worm and injected it into the Internet. He chose to release it from MIT, to disguise the fact that the worm came from Cornell. Morris soon discovered that the program was replicating and reinfesting machines at a much faster rate than he had anticipated—there was a bug. Ultimately, many machines at locations around the country either crashed or became “catatonic.” When Morris realized what was happening, he contacted a friend at Harvard to discuss a solution. Eventually, they sent an anonymous message from Harvard over the network, instructing programmers how to kill the worm and prevent reinfection. However, because the network route was clogged, this message did not get through until it was too late. Computers were affected at many sites, including universities, military sites, and medical research facilities. The estimated cost of dealing with the worm at each installation ranged from \$200 to more than \$53,000.¹

The program took advantage of a hole in the debug mode of the Unix *sendmail* program, which runs on a system and waits for other systems to connect to it and give it email, and a hole in the finger daemon *fingerd*, which serves finger requests (see Section 6.1 [Finger], page 53). People at the University of California at Berkeley and MIT had copies of the program and were actively *disassembling* it (returning the program back into its source form) to try to figure out how it worked.

Teams of programmers worked non-stop to come up with at least a temporary fix, to prevent the continued spread of the worm. After about twelve

¹ Derived in part from a letter by Severo M. Ornstein, in the *Communications of the ACM*, Vol 32 No 6, June 1989.

hours, the team at Berkeley came up with steps that would help retard the spread of the virus. Another method was also discovered at Purdue and widely published. The information didn't get out as quickly as it could have, however, since so many sites had completely disconnected themselves from the network.

After a few days, things slowly began to return to normalcy and everyone wanted to know who had done it all. Morris was later named in *The New York Times* as the author (though this hadn't yet been officially proven, there was a substantial body of evidence pointing to Morris).

Robert T. Morris was convicted of violating the computer Fraud and Abuse Act (Title 18), and sentenced to three years of probation, 400 hours of community service, a fine of \$10,050, and the costs of his supervision. His appeal, filed in December, 1990, was rejected the following March.

8.2 The Cuckoo's Egg

First in an article entitled "Stalking the Wily Hacker," and later in the book *The Cuckoo's Egg*, Clifford Stoll detailed his experiences trying to track down someone breaking into a system at Lawrence Berkeley Laboratory in California.²

A 75-cent discrepancy in the Lab's accounting records led Stoll on a chase through California, Virginia, and Europe to end up in a small apartment in Hannover, West Germany. Stoll dealt with many levels of bureaucracy and red tape, and worked with the FBI, the CIA, and the German Bundespost trying to track his hacker down.

The experiences of Stoll, and particularly his message in speaking engagements, have all pointed out the dire need for communication between parties on a network of networks. The only way everyone can peacefully co-exist in "Cyberspace is by ensuring rapid recognition of any existing problems.

8.3 Organizations

The indomitable need for humans to congregate and share their common interests is also present in the computing world. User groups exist around the world, where people share ideas and experiences. Similarly, there are organizations which are one step "above" user groups; that is to say, they

² See the bibliography for full citations.

exist to encourage or promote an idea or set of ideas, rather than support a specific computer or application of computers.

8.3.1 The Association for Computing Machinery

The Association for Computing Machinery (the ACM) was founded in 1947, immediately after Eckert and Mauchly unveiled one of the first electronic computers, the ENIAC, in 1946. Since then, the ACM has grown by leaps and bounds, becoming one of the leading educational and scientific societies in the computer industry.

The ACM's stated purposes are:

- To advance the sciences and arts of information processing;
- To promote the free interchange of information about the sciences and arts of information processing both among specialists and among the public;
- To develop and maintain the integrity and competence of individuals engaged in the practices of the sciences and arts of information processing.

Membership in the ACM has grown from seventy-eight in September, 1947, to over 77,000 today. There are local chapters around the world, and many colleges and universities endorse student chapters. Lecturers frequent these meetings, which tend to be one step above the normal "user group" gathering. A large variety of published material is also available at discounted prices for members of the association.

The ACM has a number of *Special Interest Groups* (SIGs) that concentrate on a certain area of computing, ranging from graphics to the Ada programming language to security. Each of the SIGs also publishes its own newsletter. There is a Usenet group, `comp.org.acm`, for the discussion of ACM topics. See Chapter 4 [Usenet News], page 29 for more information on reading news.

For more information and a membership application, write to:

Association for Computing Machinery
1515 Broadway
New York City, NY 10036
`ACMHELP@ACMVM.BITNET`
(212) 869-7440

8.3.2 Computer Professionals for Social Responsibility

The CPSR is an alliance of computer professionals concentrating on cer-

tain areas of the impact of computer technology on society. It traces its history to the fall of 1981, when several researchers in Palo Alto, California, organized a lunch meeting to discuss their shared concerns about the connection between computing and the nuclear arms race. Out of that meeting and the discussions which followed, CPSR was born, and has been active ever since.³

The national CPSR program focuses on the following project areas:

- **Reliability and Risk** This area reflects on the concern that overreliance on computing technology can lead to unacceptable risks to society. It includes, but isn't limited to, work in analyzing military systems such as SDI.
- **Civil Liberties and Privacy** This project is concerned with such topics as the FBI National Crime Information Center, the growing use of databases by both government and private industry, the right of access to public information, extension of First Amendment rights to electronic communication, and establishing legal protections for privacy of computerized information.
- **Computers in the Workplace** The CPSR Workplace Project has concentrated its attention on the design of software for the workplace, and particularly on the philosophy of "participatory design," in which software designers work together with users to ensure that systems meet the actual needs of that workplace.
- **The 21st Century Project** This is a coalition with other professional organizations working towards redirecting national research priorities from concentrating on military issues to anticipating and dealing with future problems as science and technology enter the next century.

For more information on the CPSR, contact them at:

Computer Professionals for Social Responsibility
P.O. Box 717
Palo Alto, CA 94302
cpsr@csli.stanford.edu
(415) 322-3778
(415) 322-3798 (Fax)

8.3.3 The Electronic Frontier Foundation

The Electronic Frontier Foundation (EFF) was established to help civilize

³ This section is part of the CPSR's letter to prospective members.

the "electronic frontier"—the Cyberspacial medium becoming ever-present in today's society; to make it truly useful and beneficial not just to a technical elite, but to everyone; and to do this in a way which is in keeping with the society's highest traditions of the free and open flow of information and communication.⁴

The mission of the EFF is

- to engage in and support educational activities which increase popular understanding of the opportunities and challenges posed by developments in computing and telecommunications;
- to develop among policy-makers a better understanding of the issues underlying free and open telecommunications, and support the creation of legal and structural approaches which will ease the assimilation of these new technologies by society;
- to raise public awareness about civil liberties issues arising from the rapid advancement in the area of new computer-based communications media and, where necessary, support litigation in the public interest to preserve, protect, and extend First Amendment rights within the realm of computing and telecommunications technology;
- to encourage and support the development of new tools which will endow non-technical users with full and easy access to computer-based telecommunications;

The Usenet newsgroups `comp.org.eff.talk` and `comp.org.eff.news` are dedicated to discussion concerning the EFF. They also have mailing list counterparts for those that don't have access to Usenet, `eff-talk-request@eff.org` and `eff-news-request@eff.org`. The first is an informal arena (aka a normal newsgroup) where anyone may voice his or her opinions. The second, `comp.org.eff.news`, is a moderated area for regular postings from the EFF in the form of *EFFector Online*. To submit a posting for the *EFFector Online*, or to get general information about the EFF, write to `eff@eff.org`. There is also a wealth of information available via anonymous FTP on `ftp.eff.org`.

The EFF can be contacted at

⁴ This section was derived from 'eff.about', available along with other material via anonymous FTP from `ftp.eff.org`

The Electronic Frontier Foundation, Inc.
155 Second St. #1
Cambridge, MA 02141
eff@eff.org
(617) 864-0665
(617) 864-0866 (Fax)

8.3.4 The Free Software Foundation

The Free Software Foundation was started by Richard Stallman (creator of the popular GNU Emacs editor). It is dedicated to eliminating restrictions on copying, redistributing, and modifying software.

The word "free" in their name does not refer to price; it refers to freedom. First, the freedom to copy a program and redistribute it to your neighbors, so that they can use it as well as you. Second, the freedom to change a program, so that you can control it instead of it controlling you; for this, the source code must be made available to you.

The Foundation works to provide these freedoms by developing free compatible replacements for proprietary software. Specifically, they are putting together a complete, integrated software system called "GNU" that is upward-compatible with Unix.⁵ When it is released, everyone will be permitted to copy it and distribute it to others. In addition, it will be distributed with source code, so you will be able to learn about operating systems by reading it, to port it to your own machine, and to exchange the changes with others.

For more information on the Free Software Foundation and the status of the GNU Project, or for a list of the current tasks that still need to be done, write to gnu@prep.ai.mit.edu.

8.3.5 The League for Programming Freedom

The League for Programming Freedom is a grass-roots organization of professors, students, businessmen, programmers and users dedicated to "bringing back" the freedom to write programs, which they contend has been lost over the past number years. The League is not opposed to the legal system that Congress intended—copyright on individual programs. Their

⁵ As an aside, the editor of the GNU project, `emacs`, contains a built-in LISP interpreter and a large part of its functionality is written in LISP. The name GNU is itself recursive (the mainstay of the LISP language); it stands for "Gnu's Not Unix."

aim is to reverse the recent changes made by judges in response to special interests, often explicitly rejecting the public interest principles of the Constitution.

The League works to abolish the new monopolies by publishing articles, talking with public officials, boycotting egregious offenders, and in the future may intervene in court cases. On May 24, 1989, the League picketed Lotus headquarters because of their lawsuits, and then again on August 2, 1990. These marches stimulated widespread media coverage for the issue. They welcome suggestions for other activities, as well as help in carrying them out.

For information on the League and how to join, write to

League for Programming Freedom
1 Kendall Square #143
P.O. Box 9171
Cambridge, MA 02139
league@prep.ai.mit.edu

8.4 Networking Initiatives

Research and development are two buzz words often heard when discussing the networking field—everything needs to go faster, over longer distances, for a lower cost. To “keep current,” one should read the various trade magazines and newspapers, or frequent the networking-oriented newsgroups of Usenet. If possible, attend trade shows and symposia like Usenix, Interop, et. al.

8.4.1 NREN

The National Research and Education Network (NREN) is a five-year project approved by Congress in the Fall of 1991. It's intended to create a national electronic “super-highway.” The NREN will be 50 times faster than the fastest available networks (at the time of this writing). Proponents of the NREN claim it will be possible to transfer the equivalent of the entire text of the Encyclopedia Britannica in one second. Further information, including the original text of the bill presented by Senator Al Gore (D-TN), is available through anonymous FTP to `nis.nsf.net`, in the directory '`nsfnet`'. In addition, Vint Cerf wrote on the then-proposed NREN in RFC-1167, *Thoughts on the National Research and Education Network*. See [RFCs], page 73 for information on obtaining RFCs.

A mailing list, 'nren-discuss@uu.psi.com', is available for discussion of the NREN; write to 'nren-discuss-request@uu.psi.com' to be added.

"To talk in publick, to think in solitude,
to read and to hear, to inquire,
and to answer inquiries, is the business of a scholar."

Samuel Johnson
Chapter VIII
The History of Rasselas, Prince of Abissinia

9 Finding Out More

9.1 Internet Resource Guide

The NSF Network Service Center (NNSC) compiles and makes available an Internet Resource Guide (IRG). The goal of the guide is to increase the visibility of various Internet resources that may help users do their work better. While not yet an exhaustive list, the guide is a useful compendium of many resources and can be a helpful reference for a new user.

Resources listed are grouped by types into sections. Current sections include descriptions of online library catalogs, data archives, online white pages directory services, networks, network information centers, and computational resources, such as supercomputers. Each entry describes the resource, identifies who can use the resource, explains how to reach the local network via the Internet, and lists contacts for more information. The list is distributed electronically by the NNSC. To receive a guide, or to get on a mailing list that alerts you to when it is updated, send a message to `resource-guide-request@nnscc.nsf.net`.

The current edition of the IRG is available via anonymous FTP from `nnscc.nsf.net`, in the directory '`/resource-guide`'.

9.2 Requests for Comments

The internal workings of the Internet are defined by a set of documents called RFCs (Request for Comments). The general process for creating an RFC is for someone wanting something formalized to write a document describing the issue and mailing it to Jon Postel (`postel@isi.edu`). He acts as a referee for the proposal. It is then commented upon by all those wishing to take part in the discussion (electronically, of course). It may go through multiple revisions. Should it be generally accepted as a good idea, it will be assigned a number and filed with the RFCs.

The RFCs can be divided into five groups: required, suggested, directional, informational and obsolete. Required RFCs (e.g., RFC-791, *The Internet Protocol*) must be implemented on any host connected to the Internet.

Suggested RFCs are generally implemented by network hosts. Lack of them does not preclude access to the Internet, but may impact its usability. RFC-793, *Transmission Control Protocol*, is a must for those implementing TCP.

Directional RFCs were discussed and agreed to, but their application has never come into wide use. This may be due to the lack of wide need for the specific application (RFC-937, *The Post Office Protocol*) or that, although technically superior, ran against other pervasive approaches (RFC-891, *Hello*). It is suggested that, should the facility be required by a particular site, an implementation be done in accordance with the RFC. This ensures that, should the idea be one whose time has come, the implementation will be in accordance with some standard and will be generally usable.

Informational RFCs contain factual information about the Internet and its operation (RFC-990, *Assigned Numbers*).

There is also a subset of RFCs called FYIs (For Your Information). They are written in a language much more informal than that used in the other, standard RFCs. Topics range from answers to common questions for new and experienced users to a suggested bibliography.

Finally, as the Internet has grown and technology has changed, some RFCs become unnecessary. These obsolete RFCs cannot be ignored, however. Frequently when a change is made to some RFC that causes a new one to obsolete others, the new RFC only contains explanations and motivations for the change. Understanding the model on which the whole facility is based may involve reading the original and subsequent RFCs on the topic.

RFCs and FYIs are available via FTP from many sources, including:

- The `nic.ddn.mil` archive, as '`/rfc/rfc-xxxx.txt`', where `xxxx` is the number of the RFC.
- from `ftp.uu.net`, in the directory '`/RFC`'.

They're also available through mail by writing to `service@nic.ddn.mil`, with a '`Subject:`' line of `send RFC-xxxx.TXT`, again with `xxxx` being the RFC number. To learn about archive servers, [Archive Servers], page 77.)

“Knowledge is of two kinds. We know a subject ourselves, or we know where we can find information upon it.”

Samuel Johnson
Letter to Lord Chesterfield
February, 1755

Conclusion

This guide is far from complete—the Internet changes on a daily (if not hourly) basis. However, this booklet should provide enough information to make the incredible breadth and complexity of the Internet a mite less imposing. Coupled with some exploration and experimentation, every user has the potential to be a competent net citizen, using the facilities that are available to their fullest.

You, the reader, are *strongly* encouraged to suggest improvements to any part of this booklet. If something was unclear, left you with doubts, or wasn't addressed, it should be fixed. If you find any problems, inaccuracies, spelling errors, etc., please report them to:

Brendan Kehoe
Department of Computer Science
Widener University
Chester, PA 19013

Internet: `guide-bugs@cs.widener.edu`
UUCP: `...!widener!guide-bugs`

If you are interested in future updates to this guide (aside from normal new editions), discussion about information to be included or removed, etc., write to '`guide-request@cs.widener.edu`' to be placed on a mailing list for such things.

“I've seed de first an de last . . . I seed de beginnin,
en now I sees de endin.”

William Faulkner
The Sound & The Fury
April 8, 1928

Appendix A Getting to Other Networks

Inter-connectivity has been and always will be one of the biggest goals in computer networking. The ultimate desire is to make it so one person can contact anyone else no matter where they are. A number of "gateways" between networks have been set up. They include:

AppleLink	Quantum Services sells access to <i>AppleLink</i> , which is similar to QuantumLink for Commodore computers and PCLink for IBM PCs and compatibles. It also provides email access through the address 'user@applelink.apple.com'.
ATTMail	AT&T sells a commercial email service called <i>ATTMail</i> . Its users can be reached by writing to 'user@attmail.com'.
BIX	Users on BIX (the Byte Information eXchange) can be reached through the DAS gateway at 'user@dcibix.das.net'.
CompuServe (CI\$)	To reach a user on the commercial service <i>CompuServe</i> , you must address the mail as <i>xxxxx.xxx@compuserve.com</i> , with <i>xxxxx.xxx</i> being their CompuServe user ID. Normally CompuServe ids are represented as being separated by a comma (like 71999,141); since most mailers don't react well to having commas in addresses, it was changed to a period. For the above address, mail would be sent to '71999.141@compuserve.com'.
EasyNet	Digital sells a service called <i>EasyNet</i> ; users that subscribe to it can be reached with the addresses <i>user@host.enet.dec.com</i> or ' <i>user%host.enet@decwrl.dec.com</i> '.
FidoNet	The FidoNet computer network can be reached by using a special addressing method. If John Smith is on the node '1:2/3.4' on FidoNet, his or her email address would be ' <i>john.smith@p4.f3.n2.z1.fidonet.org</i> ' (notice how the numbers fall in place?).
MCI Mail	MCI also sells email accounts (similar to ATTMail). Users can be reached with ' <i>user@mcimail.com</i> '.
PeaceNet	Users on the <i>PeaceNet</i> network can be reached by writing to ' <i>user@igc.org</i> '.

This table is far from complete. In addition to sites not being listed, some services are not (nor do they plan to be) accessible from the "outside" (like *Prodigy*); others, like *Genie*, are actively investigating the possibility of creating a gateway into their system. For the latest information, consult a list called the *Inter-Network Mail Guide*. It's available from a number of FTP sites, including *UUNET*; see Section 3.2.2 [Anonymous FTP], page 21, for more information on getting a copy of it using anonymous FTP.

Appendix B Retrieving Files via Email

For those who have a connection to the Internet, but cannot FTP, there do exist a few alternatives to get those files you so desperately need. When requesting files, it's imperative that you keep in mind the size of your request—odds are the other people who may be using your link won't be too receptive to sudden bursts of really heavy traffic on their normally sedate connection.

Archive Servers

An alternative to the currently well over-used FTPmail system is taking advantage of the many archive servers that are presently being maintained. These are programs that receive email messages that contain commands, and act on them. For example, sending an archive server the command 'help' will usually yield, in the form of a piece of email, information on how to use the various commands that the server has available.

One such archive server is 'service@nic.ddn.mil'. Maintained by the Network Information Center (NIC) in Chantilly, VA, the server is set up to make all of the information at the NIC available for people who don't have access to FTP. This also includes the WHOIS service (see Section 6.4.1 [Whois], page 57). Some sample 'Subject:' lines for queries to the NIC server are:

Subject: help	Describes available commands.
Subject: rfc 822	Sends a copy of RFC-822.
Subject: rfc index	Sends an index of the available RFCs.
Subject: netinfo domain-template.txt	Sends a domain application.
Subject: whois widener	Sends WHOIS information on 'widener'.

More information on using their archive server can be obtained by writing to their server address service@nic.ddn.mil with a 'Subject:' of help.

There are different "brands" of archive server, each with its own set of commands and services. Among them there often exists a common set of commands and services (e.g. 'index', 'help', etc). Be that as it may, one should always consult the individual help for a specific server before assuming the syntax—100K surprises can be hard on a system.

FTP-by-Mail Servers

Some systems offer people the ability to receive files through a mock-FTP interface via email. See Section 3.2.2 [Anonymous FTP], page 21 for

a general overview of how to FTP. The effects of providing such a service varies, although a rule of thumb is that it will probably use a substantial amount of the available resources on a system.

The "original" FTP-by-Mail service, BITFTP, is available to BITNET users from the Princeton node PUCC. It was once accessible to anyone, but had to be closed out to non-BITNET users because of the heavy load on the system.

In response to this closure, Paul Vixie designed and installed a system called FTPmail on one of Digital's gateway computers, `decwrl.dec.com`. Write to '`ftpmail@decwrl.dec.com`' with 'help' in the body of the letter for instructions on its use. The software is undergoing constant development; once it reaches a stable state, other sites will be encouraged to adopt it and provide the service also.

Appendix C Newsgroup Creation

Everyone has the opportunity to make a *Call For Votes* on the Usenet and attempt to create a newsgroup that he/she feels would be of benefit to the general readership. The rules governing newsgroup creation have evolved over the years into a generally accepted method. They only govern the "world" groups; they aren't applicable to regional or other alternative hierarchies.

Discussion

A discussion must first take place to address issues like the naming of the group, where in the group tree it should go (e.g. `rec.sports.koosh` vs `rec.games.koosh?`), and whether or not it should be created in the first place. The formal *Request For Discussion* (RFD) should be posted to `news.announce.newgroups`, along with any other groups or mailing lists at all related to the proposed topic. `news.announce.newgroups` is moderated. You should place it first in the '`Newsgroups:`' header, so that it will get mailed to the moderator *only*. The article won't be immediately posted to the other newsgroups listed; rather, it will give you the opportunity to have the moderator correct any inconsistencies or mistakes in your RFD. He or she will take care of posting it to the newsgroups you indicated. Also the '`Followup-To:`' header will be set so that the actual discussion takes place only in `news.groups`. If a user has difficulty posting to a moderated group, he or she may mail submissions intended for `news.announce.newgroups` to the address '`announce-newgroups@rpi.edu`'.

The final name and *charter* of the group, and whether it will be moderated or unmoderated, will be determined during the discussion period. If it's to be moderated, the discussion will also decide who the moderator will be. If there's no general agreement on these points among those in favor of a new group at the end of 30 days, the discussion will be taken into mail rather than continued posting to `news.groups`; that way, the proponents of the group can iron out their differences and come back with a proper proposal, and make a new *Request For Discussion*.

Voting

After the discussion period (which is mandatory), if it's been determined that a new group really is desired, a name and charter are agreed upon, and it's been determined whether the group will be moderated (and by whom), a

Call For Votes (CFV) should be posted to `news.announce.newgroups`, along with any other groups that the original Request For Discussion was posted to. The CFV should be posted (or mailed to the `news.announce.newgroups` moderator) as soon as possible after the discussion ends (to keep it fresh in everyone's mind).

The Call for Votes should include clear instructions on how to cast a vote. It's important that it be clearly explained how to both vote for and against a group (and be of equivalent difficulty or ease). If it's easier for you or your administrator, two separate addresses can be used to mail yes and no votes to, providing that they're on the same machine. Regardless of the method, everyone must have a very specific idea of how to get his/her vote counted.

The voting period can last between 21 and 31 days, no matter what the preliminary results of the vote are. A vote can't be called off simply because 400 "no" votes have come in and only two "yes" votes. The Call for Votes should include the exact date that the voting period will end—only those votes arriving on the vote-taker's machine before this date can be counted.

To keep awareness high, the CFV can be repeated during the vote, provided that it gives the same clear, unbiased instructions for casting a vote as the original; it also has to be the same proposal as was first posted. The charter can't change in mid-vote. Also, votes that're posted don't count—only those that were mailed to the vote-taker can be tallied.

Partial results should never be included; only a statement of the specific proposal, that a vote is in progress on it, and how to cast a vote. A mass acknowledgement ("Mass ACK" or "Vote ACK") is permitted; however, it must be presented in a way that gives no indication of which way a person voted. One way to avoid this is to create one large list of everyone who's voted, and sort it in alphabetical order. It should not be two sorted lists (of the yes and no votes, respectively).

Every vote is autonomous. The votes for or against one group can't be transferred to another, similar proposal. A vote can only count for the exact proposal that it was a response to. In particular, a vote for or against a newsgroup under one name can't be counted as a vote for or against another group with a different name or charter, a different moderated/unmoderated status, or, if it's moderated, a different moderator or set of moderators. Whew!

Finally, the vote has to be explicit; they should be of the form 'I vote for the group `foo.bar` as proposed' or 'I vote against the group `foo.bar` as proposed'. The wording doesn't have to be exact, your intention just has to be clear.

The Result of a Vote

At the end of the voting period, the vote-taker has to post (to `news.announce.newgroups`) the tally and email addresses of the votes received. Again, it can also be posted to any of the groups listed in the original CFV. The tally should make clear which way a person voted, so the results can be verified if it proves necessary to do so.

After the vote result is posted to `news.announce.newgroups`, there is a mandatory five-day waiting period. This affords everyone the opportunity to correct any errors or inconsistencies in the voter list or the voting procedure.

Creation of the Group

If, after the waiting period, there are no serious objections that might invalidate the vote, the vote is put to the "water test." If there were 100 more valid 'YES/create' votes than 'NO/don't' create votes, and at least two-thirds of the total number of votes are in favor of creation, then a newgroup control message can be sent out (often by the moderator of `news.announce.newgroups`). If the 100-vote margin or the two-thirds percentage isn't met, the group has failed and can't be created.

If the proposal failed, all is not lost—after a six-month waiting period (a "cooling down"), a new Request For Discussion can be posted to `news.groups`, and the whole process can start over again. If after a couple of tries it becomes obvious that the group is not wanted or needed, the vote-taker should humbly step back and accept the opinion of the majority. (As life goes, so goes Usenet.)

Glossary

This glossary is only a tiny subset of all of the various terms and other things that people regularly use on The Net. For a more complete (and very entertaining) reference, it's suggested you get a copy of *The New Hacker's Dictionary*, which is based on a VERY large text file called the Jargon File. Edited by Eric Raymond (eric@snark.thyrsus.com), it is available from the MIT Press, Cambridge, Massachusetts, 02142; its ISBN number is 0-262-68069-6. Also see RFC-1208, *A Glossary of Networking Terms*.

:-) This odd symbol is one of the ways a person can portray "mood" in the very flat medium of computers—by using "smilies." This is 'meta-communication', and there are literally hundreds of them, from the obvious to the obscure. This particular example expresses "happiness." Don't see it? Tilt your head to the left 90 degrees. Smilies are also used to denote sarcasm.

address resolution Conversion of an Internet address to the corresponding physical address. On an ethernet, resolution requires broadcasting on the local area network.

administrivia Administrative tasks, most often related to the maintenance of mailing lists, digests, news gateways, etc.

anonymous FTP Also known as "anon FTP"; a service provided to make files available to the general Internet community—see Section 3.2.2 [Anonymous FTP], page 21.

ANSI The American National Standards Institute disseminates basic standards like ASCII, and acts as the United States' delegate to the ISO. Standards can be ordered from ANSI by writing to the ANSI Sales Department, 1430 Broadway, New York, NY 10018, or by telephoning (212) 354-3300.

archie A service which provides lookups for packages in a database of the offerings of countless of anonymous FTP sites. See Section 3.3.1 [archie], page 25 for a full description.

archive server An email-based file transfer facility offered by some systems.

ARPA (Advanced Research Projects Agency) Former name of DARPA, the government agency that funded ARPAnet and later the DARPA Internet.

ARPAnet A pioneering long haul network funded by ARPA. It served as the basis for early networking research as well as a central backbone during the development of the Internet. The ARPAnet consisted of individual packet switching computers interconnected by leased lines. The ARPAnet no longer exists as a singular entity.

asynchronous Transmission by individual bytes, not related to specific timing on the transmitting end.

auto-magic Something which happens pseudo-automatically, and is usually too complex to go into any further than to say it happens "auto-magically."

backbone A high-speed connection within a network that connects shorter, usually slower circuits. Also used in reference to a system that acts as a "hub" for activity (although those are becoming much less prevalent now than they were ten years ago).

bandwidth The capacity of a medium to transmit a signal. More informally, the mythical "size" of The Net, and its ability to carry the files and messages of those that use it. Some view certain kinds of traffic (FTPing hundreds of graphics images, for example) as a "waste of bandwidth" and look down upon them.

BITNET (Because It's Time Network) An NJE-based international educational network.

bounce The return of a piece of mail because of an error in its delivery.

btw An abbreviation for "by the way."

CFV (Call For Votes) Initiates the voting period for a Usenet newsgroup. At least one (occasionally two or more) email address is customarily included as a repository for the votes. See See Appendix C [Newsgroup Creation], page 79 for a full description of the Usenet voting process.

ClariNews The fee-based Usenet newsfeed available from ClariNet Communications.

client The user of a network service; also used to describe a computer that relies upon another for some or all of its resources.

Cyberspace A term coined by William Gibson in his fantasy novel *Neuromancer* to describe the "world" of computers, and the society that gathers around them.

datagram The basic unit of information passed across the Internet. It contains a source and destination address along with data. Large messages are broken down into a sequence of IP datagrams.

disassembling Converting a binary program into human-readable machine language code.

DNS (Domain Name System) The method used to convert Internet names to their corresponding Internet numbers.

domain A part of the naming hierarchy. Syntactically, a domain name consists of a sequence of names or other words separated by dots.

dotted quad A set of four numbers connected with periods that make up an Internet address; for example, 147.31.254.130.

email The vernacular abbreviation for electronic mail.

email address The UUCP or domain-based address that a user is referred to with. For example, the author's address is `brendan@cs.widener.edu`.

ethernet A 10-million bit per second networking scheme originally developed by Xerox Corporation. Ethernet is widely used for LANs because it can network a wide variety of computers, it is not proprietary, and components are widely available from many commercial sources.

FDDI (Fiber Distributed Data Interface) An emerging standard for network technology based on fiber optics that has been established by ANSI. FDDI specifies a 100-million bit per second data rate. The access control mechanism uses token ring technology.

flame A piece of mail or a Usenet posting which is violently argumentative.

FQDN (Fully Qualified Domain Name) The FQDN is the full site name of a system, rather than just its hostname. For example, the system `lisa` at Widener University has a FQDN of `lisa.cs.widener.edu`.

FTP (File Transfer Protocol) The Internet standard high-level protocol for transferring files from one computer to another.

FYI An abbreviation for the phrase "for your information." There is also a series of RFCs put out by the Network Information Center called FYIs; they address common questions of new users and many other useful things. See [RFCs], page 73 for instructions on retrieving FYIs.

gateway A special-purpose dedicated computer that attaches to two or more networks and routes packets from one network to the other. In particular, an Internet gateway routes IP datagrams among the networks it connects. Gateways route packets to other gateways until they can be delivered to the final destination directly across one physical network.

header The portion of a packet, preceding the actual data, containing source and destination addresses and error-checking fields. Also part of a message or news article.

hostname The name given to a machine. (See also FQDN.)

IMHO (In My Humble Opinion) This usually accompanies a statement

that may bring about personal offense or strong disagreement.

Internet A concatenation of many individual TCP/IP campus, state, regional, and national networks (such as NSFnet, ARPAnet, and Milnet) into one single logical network all sharing a common addressing scheme.

Internet number The dotted-quad address used to specify a certain system. The Internet number for the site `cs.widener.edu` is 147.31.254.130. A resolver is used to translate between hostnames and Internet addresses.

interoperate The ability of multi-vendor computers to work together using a common set of protocols. With interoperability, PCs, Macs, Suns, Dec VAXen, CDC Cybers, etc, all work together allowing one host computer to communicate with and take advantage of the resources of another.

ISO (International Organization for Standardization) Coordinator of the main networking standards that are put into use today.

kernel The level of an operating system or networking system that contains the system-level commands or all of the functions hidden from the user. In a Unix system, the kernel is a program that contains the device drivers, the memory management routines, the scheduler, and system calls. This program is always running while the system is operating.

LAN (Local Area Network) Any physical network technology that operates at high speed over short distances (up to a few thousand meters).

mail gateway A machine that connects to two or more electronic mail systems (especially dissimilar mail systems on two different networks) and transfers mail messages among them.

mailing list A possibly moderated discussion group, distributed via email from a central computer maintaining the list of people involved in the discussion.

mail path A series of machine names used to direct electronic mail from one user to another.

medium The material used to support the transmission of data. This can be copper wire, coaxial cable, optical fiber, or electromagnetic wave (as in microwave).

multiplex The division of a single transmission medium into multiple logical channels supporting many simultaneous sessions. For example, one network may have simultaneous FTP, telnet, rlogin, and SMTP connections, all going at the same time.

net.citizen An inhabitant of Cyberspace. One usually tries to be a good net.citizen, lest one be flamed.

netiquette A pun on "etiquette"; proper behavior on The Net. See Section 4.13 [Usenet Netiquette], page 37.

network A group of machines connected together so they can transmit information to one another. There are two kinds of networks: local networks and remote networks.

NFS (Network File System) A method developed by Sun Microsystems to allow computers to share files across a network in a way that makes them appear as if they're "local" to the system.

NIC The Network Information Center.

node A computer that is attached to a network; also called a host.

NSFnet The national backbone network, funded by the National Science Foundation and operated by the Merit Corporation, used to interconnect regional (mid-level) networks such as WestNet to one another.

packet The unit of data sent across a packet switching network. The term is used loosely. While some Internet literature uses it to refer specifically to data sent across a physical network, other literature views the Internet as a packet switching network and describes IP datagrams as packets.

polling Connecting to another system to check for things like mail or news.

postmaster The person responsible for taking care of mail problems, answering queries about users, and other related work at a site.

protocols A formal description of message formats and the rules two computers must follow to exchange those messages. Protocols can describe low-level details of machine-to-machine interfaces (e.g., the order in which bits and bytes are sent across a wire) or high-level exchanges between allocation programs (e.g., the way in which two programs transfer a file across the Internet).

recursion The facility of a programming language to be able to call functions from within themselves.

resolve Translate an Internet name into its equivalent IP address or other DNS information.

RFD (Request For Discussion) Usually a two- to three-week period in which the particulars of newsgroup creation are battled out.

route The path that network traffic takes from its source to its destination.

router A dedicated computer (or other device) that sends packets from one place to another, paying attention to the current state of the network.

RTFM (Read The Fantastic Manual) . This anacronym is often used when someone asks a simple or common question. The word 'Fantastic' is usually replaced with one much more vulgar.

SMTP (Simple Mail Transfer Protocol) The Internet standard protocol for transferring electronic mail messages from one computer to another. SMTP specifies how two mail systems interact and the format of control messages they exchange to transfer mail.

server A computer that shares its resources, such as printers and files, with other computers on the network. An example of this is a Network File System (NFS) server which shares its disk space with other computers.

signal-to-noise ratio When used in reference to Usenet activity, 'signal-to-noise ratio' describes the relation between amount of actual information in a discussion, compared to their quantity. More often than not, there's substantial activity in a newsgroup, but a very small number of those articles actually contain anything useful.

signature The small, usually four-line message at the bottom of a piece of email or a Usenet article. In Unix, it's added by creating a file '.signature' in the user's home directory. Large signatures are a no-no.

summarize To encapsulate a number of responses into one coherent, usable message. Often done on controlled mailing lists or active newsgroups, to help reduce bandwidth.

synchronous Data communications in which transmissions are sent at a fixed rate, with the sending and receiving devices synchronized.

TCP/IP (Transmission Control Protocol/Internet Protocol) A set of protocols, resulting from ARPA efforts, used by the Internet to support services such as remote login (telnet), file transfer (FTP) and mail (SMTP).

telnet The Internet standard protocol for remote terminal connection service. Telnet allows a user at one site to interact with a remote timesharing system at another site as if the user's terminal were connected directly to the remote computer.

terminal server A small, specialized, networked computer that connects many terminals to a LAN through one network connection. Any user on the network can then connect to various network hosts.

TeX A free typesetting system by Donald Knuth.

twisted pair Cable made up of a pair of insulated copper wires wrapped around each other to cancel the effects of electrical noise.

UUCP (Unix to Unix Copy Program) A store-and-forward system, primarily for Unix systems but currently supported on other platforms (e.g. VMS and personal computers).

WAN (Wide-Area Network) A network spanning hundreds or thousands of miles.

workstation A networked personal computing device with more power

than a standard IBM PC or Macintosh. Typically, a workstation has an operating system such as unix that is capable of running several tasks at the same time. It has several megabytes of memory and a large, high-resolution display. Examples are Sun workstations and Digital DECstations.

worm A computer program which replicates itself. The Internet worm (see Section 8.1 [The Internet Worm], page 63) was perhaps the most famous; it successfully (and accidentally) duplicated itself on systems across the Internet.

wrt With respect to.

“I hate definitions.”
Benjamin Disraeli
Vivian Grey, bk. i chap ii

Bibliography

What follows is a compendium of sources that have information that will be of use to anyone reading this guide. Most of them were used in the writing of the booklet, while others are simply noted because they are a must for any good net.citizen's bookshelf.

Books

Comer, Douglas E. (1991). *Internetworking With TCP/IP, 2nd ed.*, 2v. Prentice Hall: Englewood Cliffs, NJ.

Davidson, John (1988). *An Introduction to TCP/IP*. Springer-Verlag: Berlin.

Frey, Donnalyn, and Adams, Rick (1989). *!@%:: A Directory of Electronic Mail Addressing and Networks*. O'Reilly and Associates: Newton, MA.

Gibson, William (1984). *Neuromancer*. Ace: New York, NY.

LaQuey, Tracy (1990). *Users' Directory of Computer Networks*. Digital Press: Bedford, MA.

Levy, Stephen (1984). *Hackers: Heroes of the Computer Revolution*. Anchor Press/Doubleday: Garden City, NY.

Partridge, Craig (1988). *Innovations in Internetworking*. ARTECH House: Norwood, MA.

Quarterman, John S. (1989). *The Matrix: Computer Networks and Conferencing Systems Worldwide*. Digital Press: Bedford, MA.

Raymond, Eric (ed) (1991). *The New Hacker's Dictionary*. MIT Press: Cambridge, MA.

Stoll, Clifford (1989). *The Cuckoo's Egg*. Doubleday: New York.

Tanenbaum, Andrew S. (1988). *Computer Networks, 2d ed.* Prentice-Hall: Englewood Cliffs, NJ.

Todinao, Grace (1986). *Using UUCP and USENET: A Nutshell Handbook*. O'Reilly and Associates: Newton, MA.

The Waite Group (1991). *Unix Communications, 2nd ed.*. Howard W. Sams & Company: Indianapolis.

Periodicals & Papers

Barlow, J. Coming Into The Country. *Communications of the ACM* 34:3 (March 1991): 2. Addresses "Cyberspace"—John Barlow was a co-founder of the EFF.

Collyer, G., and Spencer, H. News Need Not Be Slow. *Proceedings of the 1987 Winter USENIX Conference*: 181–90. USENIX Association, Berkeley, CA (January 1987).

Denning, P. The Internet Worm. *American Scientist* (March–April 1989): 126–128.

_____. The Science of Computing: Computer Networks. *American Scientist* (March–April 1985): 127–129.

Frey, D., and Adams, R. USENET: Death by Success? *UNIX REVIEW* (August 1987): 55–60.

Gifford, W. S. ISDN User-Network Interfaces. *IEEE Journal on Selected Areas in Communications* (May 1986): 343–348.

Ginsberg, K. Getting from Here to There. *UNIX REVIEW* (January 1986): 45.

Hiltz, S. R. The Human Element in Computerized Conferencing Systems. *Computer Networks* (December 1978): 421–428.

Horton, M. What is a Domain? *Proceedings of the Summer 1984 USENIX Conference*: 368–372. USENIX Association, Berkeley, CA (June 1984).

Jacobsen, Ole J. Information on TCP/IP. *ConneXions—The Interoperability Report* (July 1988): 14–15.

Jennings, D., et al. Computer Networking for Scientists. *Science* (28 February 1986): 943–950.

Markoff, J. "Author of computer 'virus' is son of U.S. electronic security expert." *New York Times* (Nov. 5, 1988): A1.

_____. "Computer snarl: A 'back door' ajar." *New York Times* (Nov. 7, 1988): B10.

McQuillan, J. M., and Walden, D. C. The ARPA Network Design Decisions. *Computer Networks* (1977): 243–289.

Ornstein, S. M. A letter concerning the Internet worm. *Communications of the ACM* 32:6 (June 1989).

Partridge, C. Mail Routing Using Domain Names: An Informal Tour. *Proceedings of the 1986 Summer USENIX Conference*: 366–76. USENIX Association, Berkeley, CA (June 1986).

Quarterman, J. Etiquette and Ethics. *ConneXions—The Interoperability Report* (March 1989): 12-16.

_____. Notable Computer Networks. *Communications of the ACM* 29:10 (October 1986). This was the predecessor to *The Matrix*.

Raeder, A. W., and Andrews, K. L. Searching Library Catalogs on the Internet: A Survey. *Database Searcher* 6 (September 1990): 16-31.

Seeley, D. A tour of the worm. *Proceedings of the 1989 Winter USENIX Conference*: 287-304. USENIX Association, Berkeley, CA (February 1989).

Shulman, G. Legal Research on USENET Liability Issues. *;login: The USENIX Association Newsletter* (December 1984): 11-17.

Smith, K. E-Mail to Anywhere. *PC World* (March 1988): 220-223.

Stoll, C. Stalking the Wily Hacker. *Communications of the ACM* 31:5 (May 1988): 14. This article grew into the book *The Cuckoo's Egg*.

Taylor, D. The Postman Always Rings Twice: Electronic Mail in a Highly Distributed Environment. *Proceedings of the 1988 Winter USENIX Conference*: 145-153. USENIX Association, Berkeley, CA (December 1988).

U.S. Gen'l Accounting Ofc. Computer Security: Virus Highlights Need for Improved Internet Management. *GAO / IMTEC- 89 - 57*, (1989). Addresses the Internet worm.

“And all else is literature.”

Paul Verlaine
The Sun, New York

Index

A

ACM 65
address, email 11, 14, 38
address, IP (Internet) 7, 20, 47
administrivia 15, 83
ANSI 83
AppleLink 75
archie 25, 26
archive servers 11, 72
ARL 59

B

bang path 11
BITFTP 78
BITNET 8
book bugs 73
bounce, mail delivery 14

C

CARL 46, 48
Clarinet 60
CompuServe 75
CPSR 65
crossposting 41
Cyberspace 1, 63, 64, 67, 84

D

databases 60
domains 5, 11, 14, 56, 85

E

EFF 66
EFF (Electronic Frontier Foundation) 6
Electronic Mail 11
extragalactic database 51

F

FAQs 43
FEDIX 50

finger 53, 63
FQDN 6, 7
Freenet 47
FSF (Free Software Foundation) 68
FTPable Items 16, 17, 43, 47, 57, 67, 75

G

gateway, mail-news 35, 37, 75
GNU Project 68

H

headers 13
help, with archie 26
help, with FTPmail 78
help, with geo server 50
help, with listservs 17

I

Internet number 7
Internet worm 63, 89
IRG (Internet Resource Guide) 71

J

journals 59

K

Knowbot 48

L

leased line 9
libraries 27, 46, 71
listserv 16, 17, 49, 59
LPF 68

M

mailing list 15
MCI Mail 75
minority scholarships 50

moderation, of newsgroups .. 34, 35, 79, 80
Morris, Robert (Jr.) 63

RFC-822, email format 13
RFCs (Requests for Comments) 37, 71

N

NED 51
newsgroups, for testing 39
NJE protocol, for BITNET 8
NNTP 36, 60
Nutshell Books 12

O

OCEANIC 51
octet 7

P

ping 54
postmaster 15, 16, 56, 87

Q

quotes, stock 60

R

resolving 7, 20, 86

S

security 54
signature files 37
SLIP links, modem-based IP 9, 19
STIS 51
Stoll, Cliff 64
subnet 7
Sun Managers 15

T

talk 55

U

UUCP 8, 11, 13, 36, 60, 73, 85
UUNET 20, 32

W

Weather 50
White Pages Pilot Project 48
WHOIS databases 55

NATIONAL AGRICULTURAL LIBRARY



1022401507